Consciousness is Data Compression

Phil Maguire (pmaguire@cs.nuim.ie)
Department of Computer Science, NUI Maynooth
Co.Kildare, Ireland

Rebecca Maguire (rebecca.maguire@dbs.ie)
Department of Psychology, Dublin Business School
34/35 South William Street, Dublin 2, Ireland

Abstract

In this article we advance the conjecture that conscious awareness is equivalent to data compression. Algorithmic information theory supports the assertion that all forms of understanding are contingent on compression (Chaitin, 2007). Here, we argue that the experience people refer to as consciousness is the particular form of understanding that the brain provides. We therefore propose that the degree of consciousness of a system can be measured in terms of the amount of data compression it carries out.

Keywords: Information theory, data compression, Solomonoff induction, phenomenal experience, Turing test.

Introduction

According to Einstein, the most incomprehensible thing about the world is that it is comprehensible. But what does it mean to comprehend? A common feature of understanding in both science and mathematics is that it involves the reduction of a set of observations or truths to a more basic set of assumptions. Indeed, Chaitin (2007) has proposed that all forms of understanding can be viewed as instances of data compression. Have a look at the sequence below and see if you can 'understand' it:

$$4, 6, 8, 12, 14, 18, 20, 24...$$

What is involved in understanding this sequence? Intuitively, one searches for a pattern that links all of the numbers together. If the numbers were randomly selected, then, more than likely, no pattern could be identified. In this case the sequence could not be described any more concisely: it would be incompressible. However, the above sequence seems amenable to compression. For example, one can posit the following hypothesis: "start at 4 and keep adding 2, except if the digits of the previous number sum to 2, 5 or 8, in which case add 4". These instructions provide a complete description of the sequence. However, because the description seems somewhat unwieldy, it is not particularly convincing. A more concise description is possible: "go through all odd prime numbers and add 1". Because this hypothesis is more concise, it intuitively reflects a deeper understanding of the sequence.

Scientific understanding is furthered by exposing greater levels of redundancy in observational data. The goal of the scientist is to craft a model which can describe a dataset in more concise terms. These models are called *theories*. The more compression a theory achieves, the greater its value. For example, Kepler's heliocentric model of the heavens is considered superior to Ptolemy's geocentric model, because it manages to describe astronomical observations in terms of three simple mathematical laws rather than a convoluted set of epicycles.

The idea that compression underpins scientific endeavor is not new. Occam's razor is a fundamental scientific principle which is attributed to the 14th century English friar, William of Ockham. The principle states that the explanation of any phenomenon should make as few assumptions as possible, eliminating those that make no difference to the observable predictions: "entities should not be multiplied unnecessarily". This law of parsimony implies that if you have two competing theories which both describe a phenomenon, the simpler (i.e. more compressed) explanation is better.

Algorithmic Information Theory

As homo sapien sapiens (Latin for *knowing man*), the urge to understand is a defining characteristic of our species. But why is it that we should devote so much energy to understanding the world around us? In order to answer this question we must turn to algorithmic information theory. Algorithmic information theory is a field which brings together mathematics, logic and computer science. The foundations of this field were laid by Chaitin, Solomonoff and Kolmogorov in the 1960s (see Li & Vitányi, 1997). According to Chaitin, it is "the result of putting Shannon's information theory and Turing's computability theory into a cocktail shaker and shaking vigorously". The basic idea is that the complexity of an object can be represented by the size of the smallest program for computing it. This new way of thinking about information was first proposed by Solomonoff (1964) and subsequently independently identified by Kolmogorov and Chaitin.

Algorithmic theory provides a clear answer as to why organisms should seek to compress observational data. Specifically, Solomonoff's (1964) theory of inductive inference reveals that compression is a necessary component of prediction. The theory provides a universal measure of the probability of an object by taking into account all of the ways in which it might have been produced. This universal a

priori probability can then be incorporated into Bayes' rule for inductive inference in order to make optimal predictions based on a set of prior observations.

Solomonoff's theory of inductive inference reveals that the more a set of observations can be compressed, the more accurately subsequent events can be predicted. Consider again the sequence 4, 6, 8, 12, 14, 18, 20, 24... The long-winded description predicts that the next number in the sequence will be 26, while the more succinct description predicts that 30 will follow. According to Solomonoff's theory, the latter must be the better prediction, because it involves a fewer number of assumptions: the shorter the length of the description, the more likely it is to be correct.

Algorithmic information theory reveals that compression is the *only* systematic means for generating predictions based on prior observations. All successful predictive systems, including animals and humans, are approximations of algorithmic induction. All useful contributions to human knowledge work by coaxing people into modifying their inductive strategies in such a way that they better approximate algorithmic induction.

In order to thrive in an uncertain environment, organisms must be able to anticipate future events; the more efficiently they can compress their experiences, the more accurate these predictions will be. Consequently, organisms have evolved brains which are prodigious compressors of information: compressing sensory information provides them with an 'understanding' of their environment (see Chater & Vitányi, 2002; Schmidhuber, 2006; Wolff, 1993). Tononi (2008) has proposed that the feeling of being conscious must be linked in some way to the integration of information which occurs in the brain. In the following sections we specify precisely the relationship between information processing and subjective awareness: specifically, we argue that the experience people describe as consciousness is equivalent to the compression that the brain carries out. Henceforth, this idea is referred to as the 'compression conjecture'. It should be noted that the conjecture does not merely suppose an association between consciousness and compression; rather it asserts that no meaningful distinction can be drawn between the two concepts.

Consciousness

From an evolutionary perspective, the sole purpose of the brain is to produce behavior that optimizes the reproductive success of an organism and its genetic material. Features of the brain which are not linked to optimizing behavior should therefore not have been rigorously preserved by evolution. Why then should brains go to the trouble of producing consciousness?

Algorithmic information theory tells us that the key to enhancing prediction (and hence reproductive success) is to optimize data compression. If the principal evolutionary pressure determining the structure of the brain has been its capacity to compress data, and if brains are the only system we know of that support consciousness, then this suggests a rigorous link between consciousness and compression. Systems that are good at compressing data seem to produce consciousness. But why should this be the case?

The Brain as a Compressor

In order to answer this question, we must consider the nature of the compression that the brain carries out. In other words, what type of understanding does the brain provide?

The success of an organism is dependent on cooperation between all of its constituent components. In order to achieve the goal of reproduction, it must exhibit coordinated behavior. For example, it does not make sense for an organism's legs to maintain independent agenda. Because the interests of both legs are intimately bound, it is more productive for them to cooperate with each other in achieving a single set of objectives (e.g. putting one foot forward while the other stays on the ground). Accordingly, the brain sources sensory information from all over the body and compresses it *in parallel*, thereby optimizing predictive accuracy for the organism as a whole. Tactile information from every limb is compressed in parallel with visual information from the eyes and audio information from the ears, giving rise to a form of understanding that is *centralized* and representative of the organism's experiences as a singular unit. The resulting decisions of the organism also appear centralized: to the external observer it seems as if the organism's body is being 'controlled' by a single entity with a singular set of objectives.

Not only does the success of an organism depend on cooperation between its constituent components, it also depends on cooperation between its past and future states. Snapshots of an organism's behavior taken at different points in time again reveal evidence of a singular set of objectives. For example, if you know you will be hungry in several hours time, you might pack a lunchbox in your bag. In this case, you are cooperating with your future self. From an evolutionary perspective, organisms cooperate with their future selves because reproduction is a challenging task which requires coordinated behavior manifested over an extended period of time. As a result, the brain goes to the effort of distilling memories which are maintained with the expectation that they will facilitate data compression at a future point.

The utility of memory can again be explained in terms of enhancing algorithmic induction. Memory allows us to make greater sense of the world by enhancing our ability to carry out compression. Incoming sensory data are compressed *in parallel* with stored historical data, allowing redundancy to be identified more efficiently and, consequently, enhancing predictive accuracy. Thus, the form of understanding that the brain produces unites not only distributed sensory organs but also past and current states of an organism. The compression conjecture proposes that the experience of this unitary form of understanding is what we mean when we use the term 'consciousness'.

Self-Awareness

Intuitively, the above account does not seem fully satisfactory. For example, one might conceive of an artificial compressor which compresses large amounts of current and historical data in parallel, though without experiencing the same form of awareness that we humans are familiar with. Indeed, the compression carried out by the brain has one additional ingredient which sets it apart from simpler compression systems: it compresses its observations of its own behavior. The capacity for a system to model its own actions necessarily involves the identification of itself as an entity separate to its surroundings. As a result, self-compression entails self-awareness.

The human brain is a self-representational structure which seeks to understand its own behavior. For example, people model their own selves in order to more accurately predict how they are going to feel and react in different situations. They build up internal models about who they think they are and use these models to inform their decisions. In addition, the human brain compresses the observed behavior of other organisms. When we watch other individuals, we realize that there is a great deal of redundancy in their activity: rather than simply cataloguing and memorizing every action they perform, we can instead posit the more succinct hypothesis of a concise 'self' which motivates these actions. By representing this self we can then make accurate predictions as to how the people around us will behave. The idea that the actions of an organism are controlled by a singular self is merely a theoretical model which eliminates redundancy in the observed behavior of that organism. People apply this same process to themselves: what you consider to be the essence of you is simply a model which compresses your observations of your own past behavior.

Phenomenality

A significant obstacle to providing a fully satisfactory theory of consciousness lies in explaining the phenomenon of subjective experience: why is it that we experience qualia which seem to elude scientific description? According to the consciousness conjecture, the 'flavor' of a quale can be linked to the particular form of compression that the brain carries out in response to a stimulus.

If an organism perceives a stimulus, yet can discern no pattern in the sensory data, then that stimulus will appear completely random and meaningless to the organism: the stimulus will not be experienced at all. On the other hand, if some redundancy can be identified, then the stimulus can be 'understood' (i.e. experienced) by relating it to previously gathered sensory information. For example, when people look at an apple, they perceive a round shape by identifying redundancy between the appearance of the apple and previously encountered round objects; they perceive a green color by identifying redundancy between the appearance of the apple and previously encountered green objects. When we 'see' an apple we are not just processing an instantaneous visual stimulus but, rather, compressing a set of data which has been gathered over a wide cross section of space and time. The structure of the brain allows a sensory stimulus to be translated into the subjective experience of understanding through the process of compression.

In sum, people don't passively observe the world around them; they gaze through the lens of understanding provided by their brains. When people talk about their subjective experience they are referring to the particular form of compression that their brain provides. The reason that these qualitative descriptions differ from objective scientific descriptions is because the subjective experience of a stimulus is dependent on how it is processed. The particular 'flavors' of qualia that we humans are familiar with are artifacts of our cognition, which are determined by the patterns our brains have evolved to detect and encode.

Describing Qualia

Intuitively, qualia appear to resist objective description. However, this intuition must be flawed, for if qualia could not be recorded in some informational form in the brain then we would not be able to remember them. In this case, all current subjective experiences would seem random and meaningless because there would be no previous subjective experiences with which to reconcile them.

According to the compression conjecture, which supposes that subjective experience and data compression are equivalent, it should be possible to provide a full description of a quale by detailing the compression that a system achieves in response to a stimulus. Thus, for example, the experience of red could be captured by describing the changing structure of the brain in response to the sight of a red object. This experience could then be comprehensively represented in terms of bits of bytes and could feasibly be contained in a book. Yet, intuitively, a book containing symbols could never capture the experience of the color red in the same way that we feel it; leafing through the pages of the book would not give rise to the subjective feeling of red. How can this apparent incongruity be rationalized?

The compression conjecture indicates that even if a book does carry a complete description of a subjective experience, merely reading the book is not sufficient for reproducing that experience. To appreciate it, the reader must be capable of compressing the data in the same manner in which it was originally compressed. For example, rather than simply leafing apathetically through pages of symbols, the reader must be capable of identifying the underlying patterns which link those symbols together. If a system is incapable of compressing the data, then it cannot 'understand' the experience which is contained within. Experience is dependent on the system which is doing the experiencing, as opposed to being intrinsic to a stimulus. Because reading a description of compression will not necessarily cause the same compression to occur in your own brain, reading about the experience of red will not make you experience red.

The Hard Problem

Initially, it might not be clear that the above satisfactorily addresses the hard problem of consciousness, which Chalmers (1995) identifies as the question of why consciousness *feels* like anything at all. In order to tackle this question, let us consider the case of an assembly of coordinated neurons (or, indeed, logic gates) called Amy. If we observe Amy's behavior over time, we will notice considerable redundancy in her actions. We can compress Amy's behavior through the succinct hypothesis of a core centralized self which is motivating her actions and which feels experiences. But this is just an abstract hypothesis based on a dataset: why should the formation of a hypothesis result in experience? The answer to this question lies in the realization that the hypothesis of Amy's subjective experience is a hypothesis which Amy herself holds, an understanding which is manifested through the compression she carries out. Understanding the hypothesis that one is feeling something and the actual experience of feeling are the *same thing*. Amy's feeling therefore exists relative to the assumption of her own existence, an assumption which the system itself is capable of making.

Conscious Systems

Algorithmic information theory makes clear predictions regarding what systems are conscious: objects which carry out compression are conscious, all other objects are not. Let us consider a chair. Intuitively, we would not expect a chair to be conscious. Can this intuition be justified by the compression conjecture?

Chairs do not carry out compression. They do not source sensory information from multiple locations and process it in parallel. They do not store memories to enhance future compression. And they do not develop a theory of self by compressing their own actions. Therefore they are not conscious.

Imagine holding a flame to the leg of a chair. The flame leaves a black mark, therefore the chair has certainly been affected by the flame. But intuitively, it does not seem reasonable to claim that the chair has *experienced* the flame. This difference between effect and experience is directly related to compression: specifically, the chair fails to experience the flame because the information it provides is not compressed in any way. If a chair's leg is burned it has no effect on any of the other legs. No information is communicated, and consequently there is no inter-leg data compression to bind the experiences of the chair together. Furthermore, the chair stores no memory (other than a black mark). The burning event has no effect on how subsequent events are processed, meaning that the experiences of the chair are not bound together across time. Finally, because the chair does not compress its own response to the flame, it has no awareness of any subjective experience.

In contrast, if a flame is held to the leg of a human, it has an immediate effect on how information from all other parts of the body is processed. The brain also stores a memory of being burned, thus altering the individual's future behavior in a manner which reflects the interests of the system as a whole. People 'feel' the effect of being burned because the compression carried out by their brain reflects an understanding of what it feels like to be burned. In contrast, no matter how many times you burn a chair, it will never react any differently.

Artificial Consciousness

The consciousness conjecture suggests that any system that carries out compression can be considered conscious to some extent. However, it should be noted that no known system is capable of matching or even approaching the depth of compression carried out by the organic brain.

Although computer algorithms such as Lempel-Ziv and BZip2 are used to compress files and text, these programs simply skim through data looking for trivial redundancy. Such compressors cannot realistically be described as 'understanding' text because the only patterns they can identify are based on simple statistical repetitions of symbols. In contrast, when people read a book they can 'explain' the text in terms of an underlying narrative derived from their own experiences of the world, a feat which involves a much deeper level of compression.

Nevertheless, there is no theoretical obstacle that would prevent consciousness from being implemented in an artificial medium. Any system that is arranged and updated in a way which allows for the compression of information will support consciousness, be it implemented in windmills, beer cans or toilet rolls. Although toilet rolls take up a lot more space and interact a lot more slowly, they can be arranged in such a manner so as to perfectly replicate the compression carried out by neurons in the brain.

Of course, the idea that a conscious being could be implemented in toilet rolls is very unsatisfactory. Such an implementation exacerbates the hard problem of reconciling a clearly reducible system with the feeling of intuitively irreducible experiences. One might ask: where does the consciousness reside? In this case the consciousness is not a property of any particular toilet roll. Rather, it is a property of the toilet roll system as a whole. Just like the behavior of a human, the output of the toilet roll system exhibits deep redundancy which can be effectively compressed through the hypothesis of a single centralized 'self'. In particular, the toilet roll system is itself aware of this hypothesis, and uses the theory of selfhood to guide its processing. The consciousness of the system therefore resides in its capacity to understand (i.e. compress) what it senses, thereby identifying itself as an entity separate to its environment.

The Location of Consciousness

Thus far, we have used the term 'compression' without describing precisely how compression can be identified in the brain. Where is it to be found? Intuitively, people assume that conscious experience must be drawn together at a single point, an idea which Dennett (1991) derisively refers to as the 'Cartesian theatre'. However, brain imaging studies indicate that cognitive processing is widely

distributed and does not appear to be bound at any particular point in space or time (Zeki, 2003).

Although intuition might suggest the need for a Cartesian theatre, it is important to note that the evolutionary demands which have shaped the brain's structure have not required information processing to be integrated in this way. The only moment that the brain is required to bring information together is when some action must be elicited; furthermore, only data relevant to that action needs to be integrated. Outside of this constraint, processing can remain distributed in space and time, with no impact on the success of the organism.

Accordingly, external time and 'conscious time' need not be synchronized to any greater extent other than to facilitate the undertaking of action when required. However, conscious observers have no possible means for observing any distribution in their consciousness relative to the environment: whenever they act on their surroundings the appropriate information processing is pulled together 'just in time'. Since it always appears to the observer as if they are embodied at a particular point in space and time, this leads them to mistakenly assume that their consciousness must be brought together at a single point in the brain, giving rise to the Cartesian theatre fallacy.

How Does the Brain Create Consciousness?

One of the goals of consciousness research is to identify how it is created in the brain: which neural structures support consciousness and which are merely superfluous biological apparatus? Using elementary computability theory we will prove that, if the compression conjecture holds, then the goal of identifying a complete theory of consciousness is unattainable.

Let us imagine that somebody someday submits a theory which offers a full description of how the brain produces consciousness. The theory is complete, meaning that it is capable of identifying precisely which structures in the brain give rise to consciousness, separating the conscious part from the non-conscious meat. Now, of course, the reviewers wish to check that the theory is correct. Accordingly, they apply the theory to their own brain activity to see whether the predictions match their experience. However, this raises the question: are the reviewers able to define their own consciousness, as required to validate the theory? Is it possible for a system to define its own self? In fact, computability theory rules this out, meaning that a complete theory of consciousness is not possible.

According to the compression conjecture, the recognition of one's own consciousness involves the identification of a structure which carries out the same form of compression. We can therefore present the problem formally in terms of a Turing machine which is capable of recognizing a program with the same input-output relationship. Consider a Turing machine T which takes input x and outputs 1 if $L(T) = L(x)$ (i.e. the languages recognized by T and x) and 0 if $L(T) \neq L(x)$. Is such a machine possible?

The machine T is not consistent. We can imagine another machine A which takes input x. The machine A first computes $T(A)$. If $T(A) = 1$ it then outputs $1 - T(x)$, which is the opposite of T, while if $T(A) = 0$ it then outputs $T(x)$, which is the same as T. In other words, the machine A checks to see whether T recognizes it as being equivalent or not. If T recognizes A as being equivalent then A proceeds to do the exact opposite, making it not equivalent to T. However, if T does not recognize A as being equivalent then A produces the same output at T, making it equivalent to T. There is no way around this obstacle (see Rice's theorem; Rice, 1953). Since no system can recognize an equivalent system from within itself, developing a complete theory of consciousness is not possible: the more precisely a theory attempts to define the conscious structure of the brain, the less feasible it will be to validate it.

The unrecognizability of the self has important implications for how we think about ourselves. For instance, we can never know who we really are; we can never fully explain our actions; we can never be certain as to what we are going to do next. In effect, the self is a helpless observer carried along by the compression going on in the brain. Of course, one feels like one is directing one's own actions because, as far as one is aware, one is. According to the compression conjecture, the model of the self is simply an explanatory mechanism that the brain uses to explain and predict its own behavior. As a result, the actions of the brain cannot help but be consistent with those of the self (see Gazzaniga, 1992). However, it is the activity of the brain which defines the nature of the self, rather than the other way around. Are you controlling your own actions? Certainly, but at the same time you can never know who *you* is.

Measuring Consciousness

If, as the compression conjecture supposes, consciousness is equivalent to data compression, then it should be possible to measure consciousness by quantifying the amount of compression that a system is capable of. The formal measure of compression is logical depth (see Bennett, 1988). Bennett's idea is that objects can be trivial, random or deep. Trivial objects, being completely predictable, contain no useful information; random ones, being completely unpredictable, do not contain any useful information either. In contrast, objects that are neither random nor trivial are called deep objects, because they support deep compression.

Deep objects are useful because they provide a store of mathematical work, allowing associated data to be compressed far more efficiently than can be achieved using shallower tools. Indeed, Bennett's (1988) theory implies that the concepts of 'depth' and 'intelligence' are equivalent, since the facilitation of compression that depth provides cannot be replicated by alternative means. Of all known objects, the human brain is the deepest, representing the stored mathematical work of decades of active cognitive processing on top of billions of years of evolution. The brain relies on its depth to mitigate the physical limitations on information processing imposed by its biological structure, such as limited storage capacity, processing speed and

susceptibility to degradation. The complexity of its structure allows people to effortlessly identify patterns which continue to elude the most advanced artificial intelligence programs.

The Turing Test

Turing (1950) suggested that if a computer, through a textual interface, can successfully convince a human judge that it is human, then it should be considered equal in intelligence to a human. However, the Turing test is not a reliable indicator of depth. Fooling a human judge is unlikely to require a deep program: a far simpler solution is to exploit the weaknesses of human psychology.

We propose an alternative test, involving compression, on which it is not possible to cheat. Because of its complexity, natural language provides the ideal medium for testing compressor depth. People use complex linguistic patterns to communicate with each other and assume that other speakers are capable of compressing the words they produce. If a computer system is as intelligent as a human, then it should be capable of compressing language to the same extent as a human.

According to algorithmic information theory, compression can be quantified in terms of predictive accuracy. For example, Shannon (1951) examined the human-perceived entropy of English by asking people to predict each letter in a document, one by one. The entropy rate turned out to be less than 1 bit per letter. People are able to predict language because of the fact that they 'understand' the text. In contrast, artificial compressors like BZip2 and Lempel-Ziv achieve much poorer levels of compression because they rely on predictable sequences of characters, without any regard for the deeper connections between words, sentences and narrative. If a computer was genuinely as intelligent as a human, it would be capable of matching the entropy rate of 1 bit per letter that Shannon observed.

We propose that the compression test is far more reliable and practical than the Turing test. For a start, there is no way to cheat: by definition, deep processing cannot be reproduced by any means other than underlying depth (the Slow Growth Law; see Bennett, 1988). It is also extremely quick and reliable: the probability of guessing the correct symbols decreases exponentially with the length of the test. While the Turing test is ambiguous and is affected by the gullibility of the tester, the compression test is simple, rigorous, reproducible and provides an exact measure of intelligence by means of the relative entropy score.

Conclusion

Intuitions regarding consciousness seem to create many problems which have not been satisfactorily resolved (see Dennett, 1991). In contrast, the framework we have described here can explain many of the questions regarding consciousness in an unambiguous and consistent manner.

The compression conjecture explains why a brain that evolved to optimize an organism's behavior should be associated with consciousness. It explains why consciousness is not amenable to scientific description. It explains what we mean by 'the self' and why brains provide self-awareness. It explains the apparent paradox of experiencing a singular perspective in a brain which carries out distributed processing. It predicts what systems are conscious and what systems are not; it reveals that a complete theory of consciousness is not possible. It tells us how to identify consciousness and it even provides a standard by which to measure consciousness.

The compression conjecture does not require special neuro-biological causal properties. It does not require mysterious quantum fluctuations in micro-tubules. It does not require an additional imperceptible dimension to the universe. It does not require the actions of a divine being. In fact, it requires nothing except data compression.

References

Bennett, C.H. (1988). Logical depth and physical complexity. In Herken, R., editor, *The Universal Turing Machine: A Half-Century Survey*, 227-258. Oxford: University Press.

Chaitin, G.J. (2007). The halting probability Ω: irreducible complexity in pure mathematics. *Milan Journal of Mathematics*, *75*, 291-304.

Chalmers, D.J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, *2(3)*, 200-219.

Chater, N. and Vitányi, P. (2003). Simplicity: A unifying principle in cognitive science? *Trends in Cognitive Sciences*, *7(1)*, 19-22.

Dennett, D. (1991). *Consciousness Explained*. Boston: Little, Brown and Company.

Gazzaniga, M.S. (1992). *Nature's Mind*. London: Basic Books.

Li, M. & Vitányi, P. (1997). *An Introduction to Kolmogorov Complexity*. New York: Springer.

Rice, H. G. (1953). Classes of recursively enumerable sets and their decision problems. *Transactions of the American Mathematical Society*, *74*, 358-366.

Schmidhuber, J. (2006). Developmental robotics, optimal artificial curiosity, creativity, music, and the fine arts. *Connection Science*, *18(2)*, 173-187.

Shannon, C.E. (1951). Prediction and entropy of printed English. *The Bell System Technical Journal*, *30*, 50-64.

Solomonoff, R.J. (1964). A formal theory of inductive inference. *Information and Control*, *7*, 1-22.

Tononi, G. (2008). Consciousness as integrated information: a provisional manifesto. *Biological Bulletin*, *215*, 216-242.

Turing, A.M. (1950). Computing machinery and intelligence. *Mind*, *59*, 433-460.

Wolff J.G. (1993). Computing, cognition and information compression. *AI Communications*, *6(2)*, 107-127.

Zeki, S. (2003). The disunity of consciousness. *Trends in Cognitive Sciences*, *7*, 214-218.