Durham University

Insights

Institute *of* Advanced Study

Two Concepts of Relevance and the Emergence of Mind

William Downes

Volume 9 2016 Number 7

ISSN 1756-2074

About Insights

Insights captures the ideas and work-in-progress of the Fellows of the Institute of Advanced Study at Durham University. Up to twenty distinguished and 'fast-track' Fellows reside at the IAS in any academic year. They are world-class scholars who come to Durham to participate in a variety of events around a core inter-disciplinary theme, which changes from year to year. Each theme inspires a new series of *Insights*, and these are listed in the inside back cover of each issue. These short papers take the form of thought experiments, summaries of research findings, theoretical statements, original reviews, and occasionally more fully worked treatises. Every fellow who visits the IAS is asked to write for this series. The Directors of the IAS – Veronica Strang, Rob Barton, Nicholas Saul and Martin Ward – also invite submissions from others involved in the themes, events and activities of the IAS. *Insights* is edited for the IAS by Nicholas Saul. Previous editors of *Insights* were Professor Susan Smith (2006–2009), Professor Michael O'Neill (2009–2012) and Professor Barbara Graziosi (2012–2015).

About the Institute of Advanced Study

The Institute of Advanced Study, launched in October 2006 to commemorate Durham University's 175th Anniversary, is a flagship project reaffirming the value of ideas and the public role of universities. The Institute aims to cultivate new thinking on ideas that might change the world, through unconstrained dialogue between the disciplines as well as interaction between scholars, intellectuals and public figures of world standing from a variety of backgrounds and countries. The Durham IAS is one of only a handful of comparable institutions in the world that incorporates the Sciences, Social Sciences, the Arts and the Humanities.

The focal point of the IAS is a programme of work associated with, but not exclusive to, an annual research theme. At the core of this work lies a prestigious Fellowship programme. This programme gathers together scholars, intellectuals and public figures of world standing or world-promise to address topics of major academic or public interest. Their mission is to anticipate the new and re-interpret the old, communicating across and working between disciplinary boundaries.

Every year, the Institute invites as many as twenty highly creative individuals to spend up to three months in Durham. They are located in Cosin's Hall, a magnificent and spacious 18th century mansion which, together with Durham Cathedral and Durham Castle, forms part of Palace Green, dominating the World Heritage Site of Durham Peninsula. During their stay, Fellows engage with departments and colleges, deliver public lectures and seminars, and, above all, join an international community of researchers to address the theme selected for that year. Further details of the IAS and its Fellowship programme can be found at www.durham.ac.uk/ias/fellows

Copyright

TWO CONCEPTS OF RELEVANCE AND THE EMERGENCE OF MIND

Linguistic pragmatics studies how language is used in situations. Within this field, the concept of relevance accounts for how participants communicate. Sperber and Wilson (1995) developed a cognitive theory of communication: a technical concept of relevance, which I will call 'cognitive relevance'. In this essay, I propose another mathematical concept of relevance within information theory (Cover and Thomas, 2006). This concept of 'algorithmic relevance' governs how the mind-brain works as an information processing system. But it might also explain the emergence of the higher functional levels where cognitive relevance applies. Using the two approaches to relevance, we can better relate cognitive, informational and physical descriptions of brain functioning, especially with respect to language and communication.

Cognitive Relevance

Relevance theory claims that all cognitive processing is governed by relevance. This is a property of an input stimulus to an individual. Stimuli are more or less relevant depending on the amount of positive cognitive effect they have, relative to the processing effort required to gain those effects. The relevance of an input to an individual is a trade-off between the degree of processing effort required for contextual effects and the degree of positive impact these have on the individual, improving their representation of the world. The more effects, the more relevance; the more effort, the less relevance. Therefore, the input that is *maximally relevant* to an individual is that which has the most positive cognitive effects for the least effort.

Sperber and Wilson (1995, p. 260) proposed a *cognitive principle of relevance* such that human cognition is geared to maximize relevance. This is then more specifically applied to *ostensive-inferential* communication. How does the mind-brain comprehend an utterance or action when this is intended to communicate a specific message? (I am referring here to speakers and hearers, although this explains any communicative behaviour.) The hearer's starting point is an utterance which is ostensive – it is obvious that to be relevant it must be taken as a communication. It conveys *communicative intent*. This, in turn, guarantees that the actual message, the speaker's *informative intent*, is *optimally relevant*.

The issue is how a hearer determines optimal relevance and so arrives at the speaker's informative intent. To do this, the hearer creates a context from which to draw inferences – contextual implications – about the input. This context is constructed by employing the information from their mental encyclopaedia which is most accessible. In this new context, guided by relevance, they deduce the speaker's intended message, their informative intent. This is whatever is optimally relevant, has sufficient positive cognitive effects achieved with least cognitive effort. We have a high-level cognitive theory about how the mind works.

To illustrate; a speaker says: 'Free me from all this red tape'.

Imagine a context where the speaker is a developer talking to an assistant about a planning application. The utterance is ostensive, therefore has optimal relevance and informative intent. The hearer first has to use the context to explicate the sentence uttered and grasp its content. They will probably infer that the word 'free' implies, with respect to the application, 'free from these planning regulations'. In the context, they also probably infer from 'all this red tape' that these same regulations are an unnecessary bureaucratic constraint on the speaker. That is the content of the speaker's informative intent.

Given the imperative input, 'Free me …', slightly more effort determines the speaker's propositional attitude, their desire. And, since red tape cannot be evaded, this is taken only as a wish. The hearer gains enough new information for least effort, so the speaker's intent is grasped. More information would cost more and not be worthwhile.

Why Develop this Alternative Formulation of Relevance?

This essay presents an alternative, information-based, concept of relevance. There are two motives for this. The first is straightforward. Cognitive science generally presupposes the brain is literally an information processing system. Therefore, it is useful to relate Sperber and Wilson's cognitive relevance to information processing. Sperber and Wilson themselves use the term 'information'. Informational analysis also makes relevance more explicit, because it is quantitative, and much less abstract – closer to the physical – than a purely cognitive description. Indeed, the brain events which implement information processing and from which it has emerged, are physically shaped by and model the environment. The informational approach inter-theoretically reduces Sperber and Wilson's terms 'cognitive effects' and 'processing effort' into quantities of new information and relative complexity of programs, respectively.

The second motive is to *generalize* the notion of relevance. This is because the concept becomes independent of both psychology and of semantic content or 'meaning'. *Relevance of information is a purely mathematical concept*. It applies to systems in general. We can study it in any system that has events as input and which step by step derives output to achieve the function of the system – it can be modelled by a Turing machine. We can measure the degree of relevance of input to any system. Being independent of 'meaning', systems operate 'blindly'. Yet the reformulation captures the essence of Sperber and Wilson's insight.

Since relevance governs the operation of any system that can be computationally modelled, as I believe all natural systems can be, we can explore the role of relevance in system change. And thus, we gain a new approach to *emergence*, to innovative system change.

Information

The alternative algorithmic concept of relevance is defined in terms of quantity of information, so we must sketch out what this means. The concept of information is abstracted from semantic content. It is about pure pattern, hence a mathematical concept (Shannon and Weaver, 1964). Accordingly, if we can define relevance in terms of information, it is also a mathematical concept. This is attractive because the mind-brain – a physical system – blindly gains the information it needs from its environment to serve its functioning. At this physical level of description, operations are 'mindless'. 'Mind talk' re-represents that physical phenomenon under an emergent higher-level cognitive description. A question then is how does information processing by this purely physical system 'gain' content, have human experiences, become mind-brain? My hypothesis is that algorithmic relevance within information theory provides an answer by explaining how 'higher' levels emerge from the functioning of the physical system.

What is quantity of information? It measures the *probability of an event* – a source situation – to a receiving system. The source is any situation type in which there are, minimally, two alternative possible events forming a system. A toss of a fair coin is used as an example. The possibilities at the source in this case are heads or tails, equi-probable if the coin is fair. New information is gained by a receiver if the resolution of the possibilities, whether heads or tails, is co-dependent with an equivalent change in the receiver. In the resolution of possibilities – the transition from possible to actual – the amount of information is reduced at the source and gained at the receiver. (Something has happened: not this, but this.) The amount of information gained by the receiver is the ratio between the number of contrasting possibilities existing at the source before and after the message is transmitted. Information is a *measure* of this narrowing of the possibilities at the source. The co-dependency itself is the 'channel of communication'. (This is to be distinguished from 'the medium', which is how transmission takes place.) In the simplest situation, the event of flipping a fair coin, one *bit* (binary digit) is the quantity of information transmitted; heads is 1 and tails is 0. Bits are necessary because we are dealing with two contrasting possibilities, a system of possibilities broken down into binary choices. There are more complex illustrations of information in Cherry (1966) and Dretske (1981, pp. 4f.) and yes/no possibilities are beautifully modelled by the game of 20 Questions (Cover and Thomas, 2006, p. 6).

We can make the situation more complicated by measuring the amount of information in bits with respect to successive events. Consider three tosses. Employing 0 and 1 for each toss, the number of possibilities can be represented by eight strings of three symbols. The average amount of information for each string in this equi-probable system is three bits. One can elaborate the examples using one die, where each die has six equi-probable possibilities, and then two dice, and so on, and any number of throws of the dice.

So far, the examples have involved equi-probable events. But most source situations consist of events with differing probabilities. Then, the average amount of information in the whole set differs from that of an individual string. Nevertheless, the information in an individual string is also measurable: *it is the inverse of the probability of that string's occurrence*. An improbable string conveys *more* information than a probable one. *Information is the negative log of the probability*.

Now apply this to a more complex example; a series of inputs from a source with two variable properties, fire and smoke. We code the systemic possibilities as fire (= 1) or no fire (= 0) and smoke (= 1) or no smoke (= 0). Let us assume: first, three instances; second, all the combinations are equi-probable. In that scenario, it would require 64 strings of six symbols to represent the total number of possibilities multiplied by the three instances. The average information for a string in this whole string set is six bits.

Now let us assume that some strings are *more probable* than other strings: where, in all three instances, smoke is 1 and fire is 1 and no smoke is 0 and no fire is 0 – *mutual information*. The amount of information in an individual signal can now be calculated – the inverse of its probability. Highly probable strings contain relatively less information depending on how often they have occurred out of all possible occurrences. Consider the alternative. If the *improbable* string occurred – smoke 1 and fire 0 – much more information would be gained by the receiver. I have simplified the example by leaving out the important questions of noise and redundancy. Information that the receiver already has affects the probability of the message for them and hence the quantity of new information they gain. If that is so, then it follows that the amount of new information gained by more instances of smoke 1, fire 1, is very low indeed.

Suppose that smoke 1, fire 1 is categorical, with a probability of 1; this does not mean that this situation is certain. That is just an appearance. Signals vary in probability between 1 and 0. But unless we specify a fixed number of instances, there is no limit in principle to the number of instances. Smoke 1 and fire 0 may someday contingently appear; the probability always approaches 1 but never reaches it.

What is important so far is that the *calculation of amount of information* is due to the probability patterning in data, independent of content or causation. Nevertheless, information gained about the source by the receiver is a real measurable quantity. It is no more occult than any other mathematical quantity.

Relevance as Algorithmic Complexity

Let us propose a *general principle of relevance*, modelled on Sperber and Wilson's cognitive principle of relevance; namely, that in their procedures all information processing systems are designed to maximize algorithmic relevance. *The algorithmic relevance of an input to a receiving system is the inverse of the degree of its algorithmic complexity for that system.*

The notion of *algorithmic complexity* was developed independently by Andrei Kolmogorov (1965) and Gregory Chaitin (1966). It is usually called '*Kolmogorov Complexity*', henceforth K-complexity (Aaronson, 2011; Gleick, 2011; Hunter, 2007; Cover and Thomas, 2006). The degree of algorithmic complexity of a string of data is a measure of the *length* of a description, bit by bit, which a universal Turing machine requires to specify that input. Cover and Thomas (2006, p. 3) write, 'the complexity of a string of data can be defined by the length of the shortest binary computer program for computing the string. The complexity is the minimal description length. This definition of complexity turns out to be universal; that is, computer independent, and of fundamental importance'.

K-complexity also equals Shannon information: the longer the description, the greater the input's K-complexity and the quantity of information. A description that was of equal length to the input would be maximally complex. It would also contain maximal information. Such input would be random, without pattern. Descriptions shorter than the input are shorter because they contain rules that capture generalizations about the data. Because of this, patterns enable *data compression* which reduces the length of the program. This reflects the fact that the string is less complex and contains less information. Since degree of relevance is the inverse of K-complexity, for any program, maximally relevant inputs are those which are derived by the shortest description and contain the least information. But this particular information is that which is most relevant to the program. It is exactly what is *minimally* needed for what the program is for – the function it performs in a containing system.

One important class of pattern involves properties which occur equi-probably and therefore always convey exactly the same amount of information. *When this informational relationship is a temporally ordered, law-like regularity, e.g. fire and smoke, it is categorized as causal, however that is interpreted*. True causal relationships are most effective for prediction. But any pattern enables descriptive compression. For example, consider an input string of English letters, CHAIRMITTE. An algorithm designed to specify a procedure to rearrange these letters to form any English word, with no other information available to the program, would be longer than one which utilized the prior information that the set of letters can be rearranged to form 'arithmetic' (adapted from Goldreich and Wigderson, 2008, p. 580).

The regular fire-smoke pattern in our example must make a program truly describing it shorter than a mere list. Because of this, our data is less K-complex. Since the degree of algorithmic complexity measures the information in the input, the program does so too: the longer the description, the more information; the shorter the description, the less information. This can be summarized by saying that the information in a binary string is the length of the shortest program that specifies it on a universal Turing machine.

Complexity in Context

Our next step is to relativize the K-complexity of an input to specific contexts. A string of data is K-complex for a *particular* program. Different programs embody differing information about input data depending on what the program is for. They embody exactly the information that is maximally relevant – what is minimally needed for the input to fulfil whatever function is served by the program within a containing system. Let us define this as a *context*. A context is a program which serves a function within a containing system. (Let us call a context, a *C-program*.) The K-complexity and hence amount of information in a string therefore can be measured for each C-program. It is what is most relevant with respect to its function in context.

C-programs are compressed according to the functionally efficacious patterns in the data they embody, only some of which are causal law-like regularities, e.g. the smoke-fire system. C-programs can be 'hard wired' as the result of natural or artificial selection. Alternatively, a receiving C-program could learn from new input and revise patterns of data already accessible to it in a memory, and henceforth process new data accordingly. It could revise expected probabilities of the input and thus amount of information gained. In artificial systems, the C-program is the software. *The key point is that input data's K-complexity is relative to C-programs*. Whatever its origin, whether a true general law or a heuristic short cut, compression determines length, bit by bit. The shorter the description, the more relevant that input is to that C-program's function. Relevance can be mathematically stated as a precise quantity.

The degree of relevance of a specific input string for a receiving system is inversely proportional to its algorithmic complexity for that system's C-program.

How Selection Works

This principle offers a general theory of selection. *Faced with alternatives, systems select the most relevant option in a context/C-program*. This is the option that can be derived with the least algorithmic complexity by the program, thus with minimum information, just that information necessary and sufficient to serve some function. This is a function, by accident or design, efficacious for the system, something that sustains it in being or facilitates its better operation or adaptation. Darwin's natural selection is a special case.

It is this efficacious function of the results of blind selection that gives adaptive evolution its teleological character. Any selection *must* have a function, a *telos*, that of the context/Cprogram which performs it. This efficacious function is what motivates the C-program; its accidental but functional employment of the law-like regularities that, of the available alternatives, best enables compression. We see how and why law-like regularity blindly generates new functionality or innovative system change. 'Adaptation' is synonymous with 'relevant'.

Sources of Variability

Selection falls into types according to the sources of variability. First, there may be variation either in the available input or between members of a population of systems. This changes what is most relevant. Given such options, some systems will be able to process in a more efficacious way, generating a new function or improving the performance of an old function. Second, there can be variation in context/C-program. Random changes in DNA provide one example. Furthermore, for complex learning systems with rich, accessible databases or memories, variation is internally available because the system's context/C-program can run using different internal data; in effect a choice between different contexts/C-programs. For example, a chess-playing program can recognize a new board input as a situation it has encountered before and for which it can access what previously proved the best next move. From variation, systems automatically choose the most relevant. This has nothing to do with content, only with quantity of information. This is a purely mathematical account of self-organization applicable to any system.

How Emergence Works

Such dynamic selective scenarios lead to emergence. From the range of available variables from whatever source, the system must select the most relevant leading to phenotypical functional success. Because of this, the option chosen, the innovative context/C-program becomes newly established. This excludes alternatives within the system and the population. Although the outcome is determined, to the degree that variability is random, it is unpredictable. Although each selection is based on a minimal number of steps and least quantity of information, the population of system of systems inevitably develops towards increasing structural and behavioural complexity. With this increase *emerge* new system properties, unpredictably. These have facilitated the functional efficacy, or are by-products of it. Newly emergent regular properties in turn provide new opportunities for future compression.

The increasingly complex modular system of systems is organized hierarchically. As each new level emerges containing already established regularities at a now subordinate level, these lower sub-systems perform efficacious innovative functions for the whole. This improves its functioning in its environment. The whole process is accidental and opportunistic. The grammatical and lexical features of language are examples. In general, communicators could refer to more complex situations and thus function better when they produce sentences rather than simply phrases or words – though phrases and words have constitutive functional roles in sentences.

Data ultimately consists of unique events. Categorization is the result of which sets of events, opposed to all other events, serve the emergent function. The categorization of lower levels of data at successively higher levels is a universal of complex systems. From an analytic point of view, taxonomic hierarchy is the inevitable result of successive re-representations of one category by another to serve various functions that can only be accomplished at one level. Relations between these categories automatically lead to the *deductive* relations – probabilistic or not – that characterize the explicit steps of programs. In this way, general categories minimize steps, compressing quantity of information. So categorization is a form of compression. This is illustrated by our fire-smoke example. The recognition of events as members of two distinct lower level categories is presupposed by the higher level categorization which correlates the two categories in the predictive hypothesis; if smoke, then fire. Being virtually categorical, this has a very low and therefore highly relevant quantity of information. It is easy to see how in the context of complex mobile systems it would be functional to gain this information. It could become ‘instinctive’, part of a program, or learned, stored in memory.

Emergence of Mind

Innovative system change or emergence is a result of selection governed by relevance as algorithmic complexity. This can be applied to how the mind naturally emerges from the brain, where a brain is a physical object and the mind is its functioning as an information processing system with respect to its environment.

At the brute physical level of physics, I assume that the brain is a patterning of events which occur with given probabilities, so contain potential information. At a deeper physical level this emerges from the brain taken as an entropic, thermodynamic system. If for the most part brain processing is not digital in its electrochemical structure, then at the informational level the brain is an *analog computer*. This is a physical system which solves mathematical problems where the input is a flow of continuous variable data. *Any* process of this kind can be digitized and so measured as information. Given this, brains are physical systems which process information in the technical Shannon-Kolmogorov sense. This information has, accidentally and blindly, come to enable the whole organism to function optimally with respect to its environment. The physical analog and informational digital descriptions are related without inconsistency.

I have assumed so far that lower level Shannon-Kolmogorov information processing is independent of content. However, implicit content does arise whenever input data serves a function. This idea of ‘*primitive intentionality*’ is suggested in Dretske (1980). A receiving system with a C-program can be said to be ‘about’ the pattern in the data which it selects. An observer can describe the relation using intentional language – *information gains genuine intentionality*. For example, the light sensitive cells that control stoma in photosynthesis can make a mistake faced with artificial light. An observer can describe this as a false belief, generating an ‘opaque context’ – ‘the leaf believes that the sun has risen’. This is the reason why it has opened the stoma. Although ‘primitive’, this behaviour appears intelligent. The new information gained is relevant to the leaf because it contains the minimal information necessary for its functioning. We can describe this as a ‘representation’.

The same idea lies behind Peirce’s *indexical signs* or Grice’s *natural meaning* (Jacob, 2014). The J. R. Firth-Malinowski dictum that *meaning is function in context* expresses a related notion. As physical systems gain intentionality and can be assigned ‘explicit’ truth-conditional content, we can talk of the emergence of ‘mind’. *This clearly is a matter of degree*. Finally, when systems can publicly re-represent their own processing to the degree that is useful, as our human observer has done, then we can talk of full human intentionality. This constitutes ‘mind’, as Brentano claimed (1874/1995). Our theory does not engage with the ‘hard problem’ of consciousness.

Evidence from Cognition

The above theory is remarkably consistent with the work of Andy Clark (1997, 2013; see ‘Andy Clark’ *Wikipedia: The Free Encyclopedia*). I will mention three of his themes.

First, for Clark the mind-brain is best thought of as a *Baysian predictive machine*. Such machines use probabilistic models to process their input using 'best guesses'. These are quantities assigned to the probability of their current state of knowledge. This 'best guess' processing proceeds hierarchically. Given our view, probability is *also* how information is measured, so the amount of information gained from an input equals its probability: very improbable, more information; very probable, less information. Thus, the predictions of the most probable categorization of input must also be the least algorithmically complex, hence the most relevant to the mind-brain, given its current context/C-program. In Clark's terms, this is the minimum information needed for action. Clark's minimum is our most relevant in context, a quantity of information. This is solely a matter of brain function, not a prediction with cognitive content. (The iterative nature of the predictive processing and the successive assignments may reflect the successive 'compressed; re-representations of the system, from bottom-level sensory data, to top-level functionally-specialized cortical activity.)

Second, processing is *fast and dirty*, employing the least amount of information – the most relevant – required for action. The function of a system is ultimately to enable maximally relevant action appropriate for each unique context. The role of compression deriving from a best approximation to data, (even if in error as in stereotyping) is a by-product of this fast and dirty, minimum necessary approach. *All that really matters is the function, not the heuristics used to achieve it*.

As we saw, categorization involves compression (for categorization in cognition and language see Rosch, 1973, 1978; Estes, 1994; Taylor, 2003). Categories are sets distinguished on the basis of properties. The most successful cognitive approach is *prototype theory*, pioneered by Eleanor Rosch. Prototypes provide criteria for categories. They isolate a bundle of properties of an exemplary subset which establishes a norm for a category, a larger set with fuzzy boundaries. The 'prototypical' type serves as a reference for determining a member of the category. Thus, in the larger superordinate set of all furniture, chairs are prototypical at what is called 'the basic level'; in the case of birds, American robins – much like English blackbirds – provide the prototype. Less typical subordinate categories are more fine-grained: 'bed', 'sofa bed' or 'coffee table'; 'ostrich', 'goose' or 'corvid'. 'Sheraton chair', 'Canada goose' and 'European jay' make even finer distinctions. Basic level prototypes are psychologically real. They also reflect correlational structures in the world, just to the *rough and ready degree* necessary to distinguish the category from other categories with respect to their use for *action*, as Clark emphasizes. Prototypes compress the amount of information required to distinguish a set into the minimum, reducing the complexity of the environment into just what is needed in each context.

Rosch's theory of how prototypes achieve compression is summarized by Taylor (2003, p. 52). He writes:

> Rosch argues that it is the basic level categories that most fully exploit the real-world correlation of attributes. Basic level terms cut up reality into maximally informative categories. The basic level therefore is the level in a categorization hierarchy at which the 'best' categories emerge. To do this, they '(a) maximize the number of attributes shared by members of the category; and (b) minimize the number of attributes shared with the members of other categories'.

In most contexts, the shortest C-program uses the basic level term. This 'best' *compresses* the environment into the least information to achieve the function; that is, the shortest, least K-complex, C-program.

A third theme explored by Clark is the *extended mind hypothesis* (Clark and Chalmers, 1998). This is the question, 'Where does the mind stop and the rest of the world begin?' The hypothesis is that mind is not co-extensive with any one brain. If mental activity is functionally defined, how these functions are achieved can involve systems which are 'outside' the brain. For example, if a *mental function* is most relevantly achieved using a technology, the mind extends beyond the brain. Examples abound; echo-location by radar, running a simulation on a computer, checklists and scripts, calculators and slide rules, etc. Cognitive extension also occurs when a mental function is best achieved by cooperating brains; by social activities, including communication. If brains are literally information processing systems as described above, and mind equals brain functioning, then the use of technologies and inter-system communication to achieve these brain functions means that the extended mind simply follows as a consequence.

Language and Communication

Successive re-representation, data compression, predictive, fast and dirty processing and mind extension are all exemplified by language. We distinguish between language in the broad sense – which includes semantic content and its pragmatic and sociolinguistic uses – and language in the narrow sense, the patterns that connect physical sound and conceptual meaning. Language re-representations must interface logical-conceptual formats with the articulatory and acoustic formats that control the specialized activity of speech. That is, it is logically necessary that the linguistic system must interface with meaning and sound, levels of logical form and phonological form respectively. Language connects sound and meaning.

Phonology wonderfully illustrates data compression. Speech sound is continuous. *Physically there is no division into segmental sounds or words*: no, 'f – r – e – e … m – e … f – r – o – m'. This continuous stream is re-represented in terms of a small number of classes of sounds called *phonemes* that function to distinguish words from one another. The phoneme represents compressed information about exactly what is relevant to tell words apart in context (actually morphemes, minimal units of meaning), telling specialized arrays of muscles to produce the connected speech from which exactly that phonemic information can be extracted. Even more significant, phonemes physically consist of organized bundles of acoustic – articulatory properties called '*distinctive features*' drawn from a small universal inventory (Jakobson et al., 1969). *The key point is that all the information inherent in any speech act, whatever its meaning, is ultimately compressed, coded, into a very small universal rule-governed set of distinctive acoustic – articulatory properties*. The phonological interface has the task of decompressing this code to determine the intended morpheme as it functions grammatically or lexically. This provides input to computations that connect it to a logical-conceptual structure. Because it is rule-governed, it works predictively.

Some morphemes, such as gender, case or number inflections, only code relevant grammatical information to serve as input to syntax, rules of sentence formation. Other 'lexical' morphemes, called 'lexical items', re-represent concepts, some of which are prototypes. As we saw, all classification serves a purpose. Categories like prototypes compress information into maximally useful clues that will need unpacking in communication. For this to happen, these must be mapped onto phonology or gesture and made public.

The job of syntax is to do this mapping of concepts onto sound and put the concepts together into thoughts, so that ultimately, in communication, the unique states of affairs that speakers intend to specify can be reconstructed by hearers. To do this, the rules of syntax categories recursively generate structures which are independent of specific content. *They can be used to convey any content. This enables speakers to say anything over an infinite range. Participants are only limited by their concepts and thoughts*. All the information that sentences contain is compressed into a very few general syntactico-semantic patterns which can be used to predict only *who did what to whom, where, why and when, relative to any arbitrary moment of speech and the speaker's attitude to all this* (e.g. 'Free me of all this red tape'). Utterances, speakers in unique times and situations, make this compressed lexical content (e.g. a prototype or event schema) and syntactico-semantic information publicly available for pragmatics.

The information available to pragmatics is only such compressed clues. The new information gained about the speaker's intention is only with respect to the hearer's own expectations of what has the highest probability. This is because the hearer assumes it is whatever is the most relevant with respect to the context/C-program that is the shortest, the least K-complex, that predicts it. This contains the least information because it is the most probable. *That* has to be what the speaker intended. This is guaranteed by the general principle of relevance. If they intend to communicate at all, the speaker and hearer logically *must* assume they are using the same program – live in the same world. This is called the 'principle of charity'.

The input is very compressed indeed – in fact, it is merely bundles of distinctive features. It is pure pattern. The information gained is not 'decompressed' from these patterns but reproduced anew from them. They are what are minimally necessary to do this – fast and dirty. The receiving system can assume that the compressed clues are just what is required to discover the source's systems intentions; the function of the utterance whatever it is. The general principle of relevance guarantees this. We have reinterpreted human cognitive relevance in terms of algorithmic relevance. Human communication is a special case of how *any* two or more systems in nature function to communicate the information they need.

Four Speculations

I will conclude with four brief suggestions concerning how the informational theory of relevance might figure in the further investigation of language, mind and beyond. The first suggestion involves analogy. It has long been believed that analogy is at the heart of linguistic innovation, both lexical and semantic and syntactic. In *cognitive linguistics*, metaphor, based on underlying analogies, is the fundamental structuring principle of thought, vocabulary and syntactico-semantics (Johnson, 1987; Lakoff and Johnson, 1980; Lakoff, 1987; Langacker, 1987, 1991). My suggestion is that analogy – that some target is structurally isomorphic with some source with respect to selected properties – is actually a form of information compression. It is the most relevant way to characterize the properties of the target. It vaguely delineates the minimum information about the target necessary to achieve the function of characterizing it. It says, look at the source in this way, as the most relevant model for information about the target.

The second suggestion concerns Chomsky's (1995) *minimalist program* in linguistics. This assumes, as a guide for theory construction, that the computational processes which function to relate the logical and phonological interfaces are simple and optimal: 'in the sense that particular phenomena are not overdetermined by linguistic principles and that the linguistic system is subject to economy restrictions with a least effort flavour' (Hornstein et al., p. 14; see also Boeckx, 2006, pp. 2–5). It is possible that minimalist assumptions fall out from a computational principle that all information processing is formally governed by maximal relevance.

The third suggestion concerns consciousness, the 'hard problem' mentioned above. Cognitive science explains mind in terms of sub-personal processes generally inaccessible to phenomenological scrutiny; e.g. language or visual processing. By contrast, the phenomenological level is how some of this information processing is finally re-represented as relevant 'conscious experience'; e.g. speech or visual perception. This emerges through selected re-representation within information processing by the brain as an analog computer. (Chalmers, 2010, pp. 25f, also proposes a 'dual-aspect information theory'.)

Famously, there is 'something it is like' to have mentality (Nagel, 1974). My suggestion is that conscious experience is uniquely a *biological* brain's information processing at its 'highest' level of *both* internal and extended re-representations. This level constitutes the inner and outer social worlds of subjective experience as representations embedded within a self-representation; which is a proprietary database with its own program. Certainly, consciousness is *systemically functional in multiple ways* (for summary, see Van Gulick, 2014, pp. 18–22). Most importantly, there is ongoing inner dialogue with oneself and outer dialogue with others within a society of mind-brains which together construct culture. Relevance guarantees that both these dialogues constitute functional contexts for the conscious 'workspace' (Baars, 1988).

The fourth suggestion goes beyond language and mind. It focuses on the fact that the informational theory of relevance as algorithmic complexity, because it is abstract, may be of more general philosophical and scientific interest. Aaronson (2011) makes a similar suggestion for computational complexity. My main new hypothesis is that relevance is the key mathematical property of all systems, because all systems, not just brains, can be construed as physically processing information, as analog computers. The general principle of relevance makes them dynamic. It motivates their functionality and explains their emergence and change. Nature wholly consists of such systems.

Reference List

Aaronson, S. (2011) Why philosophers should care about computational complexity. *PhilSci-Archive*.

Baars, B. (1988) *A Cognitive Theory of Consciousness*. Cambridge: Cambridge University Press.

Boeckx, C. (2006) Linguistic Minimalism: Origins, Concepts, Methods and Aims. Oxford: Oxford University Press.

Brentano, F. (1874/1995) *Psychology from an Empirical Standpoint*. Kraus, O. (ed.) Translated by Rancurello, A., Terrell, D. and McAlister, L. London: Routledge.

Chaitin, G. (1966) On the length of programs for computing finite binary sequences. *Journal of the ACM* 13: 547–69.

Chalmers, D. (2010) *The Character of Consciousness*. Oxford: Oxford University Press.

Cherry, C. (1966) *On Human Communication*. Second Edition. Cambridge, MA: MIT Press.

Chomsky, N. (1995) A Minimalist Program for Linguistic Theory. In *The Minimalist Program*. Cambridge, MA: MIT Press, pp. 167–217.

Clark, A. (1997) *Being There*. Cambridge, MA: A Bradford Book, MIT Press.

- - - . (2013) Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences* 36(3): 1–73.

Clark, A. and Chalmers, D. (1998) The extended mind. *Analysis* 58: 10–23.

Cover, T. and Thomas, J. (2006) *Elements of Information Theory*. Second Edition. New York: John Wiley and Sons.

Dretske, F. (1980) The Intentionality of Cognitive States. In French, P., Uehling, T. and Wettstein, K. (eds.) *Midwest Studies in Philosophy*, Vol. 5. Minneapolis: University of Minnesota Press, pp. 281–94.

- - - . (1981) *Knowledge and the Flow of Information*. Oxford: Blackwell.

Estes, W. (1994) *Classification and Cognition.* New York, Oxford: Oxford University Press.

Gleick, J. (2011) *The Information*. London: Fourth Estate.

Goldreich, O. and Wigderson, A. (2008) Computational Complexity. In Gowers, T. *The Princeton Companion to Mathematics*. Princeton, NJ: Princeton University Press, pp. 575–604.

Hornstein, N., Nunes, J. and Grohmann, K. (2005) *Understanding Minimalism*. Cambridge: Cambridge University Press.

Hunter, M. (2007) Algorithmic information theory. *Scholarpedia* 2(3): 2519.

Jacob, P. (2014) Intentionality. *Stanford Encyclopedia of Philosophy* (Winter). Available at http://plato.stanford.edu/archives/win2014/entries/intentionality/

Jakobson, R., Fant, C. and Halle, M. (1969) *Preliminaries to Speech Analysis*. Cambridge, MA: MIT Press.

Johnson, M. (1987) *The Body in the Mind*. Chicago: University of Chicago Press.

Kolmogorov, A. (1965) Three approaches to the quantitative definition of information. *Problems of Information and Transmission* 1(1): 3–11.

Lakoff, G. (1987) *Women, Fire and Dangerous Things*. Chicago, London: University of Chicago Press.

Lakoff, G. and Johnson, M. (1980) *Metaphors We Live By*. Chicago, London: University of Chicago Press.

Langacker, R. (1987) *Foundations of Cognitive Grammar*, Vol. 1. Stanford, CA: Stanford University Press.

- - - . (1991) *Foundations of Cognitive Grammar*, Vol. 2. Stanford, CA: Stanford University Press.

Nagel, T. (1974) What is it like to be a bat? *Philosophical Review* 83(4): 435–50. Reprinted in Hofstadter, D. and Dennett, D. (eds.) *The Mind's Eye*. London: Penguin: 391–403.

Rosch, E. (1973) Natural categories. *Cognitive Psychology* 4: 328–50.

- - - . (1978) Principles of Categorization. In Rosch, E. and Lloyd, B. (eds.) *Cognition and Categorization*. Hillsdale, NJ: Laurence Erlbaum, pp. 27–48.

Shannon, C. and Weaver, W. (1964) *The Mathematical Theory of Communication*. Urbana, IL: University of Illinois Press.

Sperber, D. and Wilson, D. (1995) *Relevance: Communication and Cognition*: Oxford: Blackwell.

Taylor, J. (2003) *Linguistic Categorization*. Third edition. Oxford: Oxford University Press.

Van Gulick, R. (2014) Consciousness. *Stanford Encyclopedia of Philosophy* (Spring). Available at http://plato.stanford.edu/archives/spr2014/entries/consciousness/

Backlist of Papers Published in Insights

No.	Author	Title	Series
2008 Volume 1			
1	Boris Wiseman	Lévi-Strauss, Caduveo Body Painting and the Readymade: Thinking Borderlines	General
2	John Hedley Brooke	Can Scientific Discovery be a Religious Experience?	Darwin's Legacy
3	Bryan R. Cullen	Rapid and Ongoing Darwinian Selection of the Human Genome	Darwin's Legacy
4	Penelope Deutscher	Women, Animality, Immunity – and the Slave of the Slave	Darwin's Legacy
5	Martin Harwit	The Growth of Astrophysical Understanding	Modelling
6	Donald MacKenzie	Making Things the Same: Gases, Emission Rights and the Politics of Carbon Markets	Modelling
7	Lorraine Code	Thinking Ecologically about Biology	Darwin's Legacy
8	Eric Winsberg	A Function for Fictions: Expanding the Scope of Science	Modelling
9	Willard Bohn	Visual Poetry in France after Apollinaire	Modelling
10	Robert A. Skipper Jr	R. A. Fisher and the Origins of Random Drift	Darwin's Legacy
11	Nancy Cartwright	Models: Parables v Fables	Modelling
12	Atholl Anderson	Problems of the 'Traditionalist' Model of Long-Distance Polynesian Voyaging	Modelling
2009 Volume 2			
1	Robert A. Walker	Where Species Begin: Structure, Organization and Stability in Biological Membranes and Model Membrane Systems	Darwin's Legacy
2	Michael Pryke	'What is Going On?' Seeking Visual Cues Amongst the Flows of Global Finance	Modelling
3	Ronaldo I. Borja	Landslides and Debris Flow Induced by Rainfall	Modelling
4	Roland Fletcher	Low-Density, Agrarian-Based Urbanism: A Comparitive View	Modelling
5	Paul Ormerod	21st Century Economics	Modelling
6	Peter C. Matthews	Guiding the Engineering Process: Path of Least Resistance versus Creative Fiction	Modelling
7	Bernd Goebel	Anselm's Theory of Universals Reconsidered	Modelling
8	Roger Smith	Locating History in the Human Sciences	Being Human
9	Sonia Kruks	Why Do We Humans Seek Revenge and Should We?	Being Human
10	Mark Turner	Thinking With Feeling	Being Human
11	Christa Davis Acampora	Agonistic Politics and the War on Terror	Being Human
12	Arun Saldanha	So What *Is* Race?	Being Human
13	Daniel Beunza and David Stark	Devices For Doubt: Models and Reflexivity in Merger Arbitage	Modelling
14	Robert Hariman	Democratic Stupidity	Being Human
2010 Volume 3			
1	John Haslett and Peter Challenor	Palaeoclimate Histories	Modelling
2	Zoltán Kövecses	Metaphorical Creativity in Discourse	Modelling
3	Maxine Sheets-Johnstone	Strangers, Trust, and Religion: On the Vulnerability of Being Alive	Darwin's Legacy

No.	Author	Title	Series
4	Jill Gordon	On Being Human in Medicine	Being Human
5	Eduardo Mendieta	Political Bestiary: On the Uses of Violence	Being Human
6	Charles Fernyhough	What is it Like to Be a Small Child?	Being Human
7	Maren Stange	Photography and the End of Segregation	Being Human
8	Andy Baker	Water Colour: Processes Affecting Riverine Organic Carbon Concentration	Water
9	Iain Chambers	Maritime Criticism and Lessons from the Sea	Water
10	Christer Bruun	Imperial Power, Legislation, and Water Management in the Roman Empire	Water
11	Chris Brooks	Being Human, Human Rights and Modernity	Being Human
12	Ingo Gildenhard and Andrew Zissos	Metamorphosis - Angles of Approach	Being Human
13	Ezio Todini	A Model for Developing Integrated and Sustainable Energy and Water Resources Strategies	Water
14	Veronica Strang	Water, Culture and Power: Anthropological Perspectives from 'Down Under'	Water
15	Richard Arculus	Water and Volcanism	Water
16	Marilyn Strathern	A Tale of Two Letters: Reflections on Knowledge Conversions	Water
17	Paul Langley	Cause, Condition, Cure: Liquidity in the Global Financial Crisis, 2007–8	Water
18	Stefan Helmreich	Waves	Water
19	Jennifer Terry	The Work of Cultural Memory: Imagining Atlantic Passages in the Literature of the Black Diaspora	Water
20	Monica M. Grady	Does Life on Earth Imply Life on Mars?	Water
21	Ian Wright	Water Worlds	Water
22	Shlomi Dinar, Olivia Odom, Amy McNally, Brian Blankespoor and Pradeep Kurukulasuriya	Climate Change and State Grievances: The Water Resiliency of International River Treaties to Increased Water Variability	Water
23	Robin Findlay Hendry	Science and Everyday Life: Water vs H_2O	Water

2011 Volume 4

No.	Author	Title	Series
1	Stewart Clegg	The Futures of Bureaucracy?	Futures
2	Henrietta Mondry	Genetic Wars: The Future in Eurasianist Fiction of Aleksandr Prokhanov	Futures
3	Barbara Graziosi	The Iliad: Configurations of the Future	Futures
4	Jonathon Porritt	Scarcity and Sustainability in Utopia	Futures
5	Andrew Crumey	Can Novelists Predict the Future?	Futures
6	Russell Jacoby	The Future of Utopia	Futures
7	Frances Bartkowski	All That is Plastic... Patricia Piccinini's Kinship Network	Being Human
8	Mary Carruthers	The Mosque That Wasn't: A Study in Social Memory Making	Futures
9	Andrew Pickering	Ontological Politics: Realism and Agency in Science, Technology and Art	Futures
10	Kathryn Banks	Prophecy and Literature	Futures
11	Barbara Adam	Towards a Twenty-First-Century Sociological Engagement with the Future	Futures
12	Andrew Crumey and Mikhail Epstein	A Dialogue on Creative Thinking and the Future of the Humanities	Futures
13	Mikhail Epstein	On the Future of the Humanities	Futures

No.	Author	Title	Series
2012 Volume 5			
1	Elizabeth Archibald	Bathing, Beauty and Christianity in the Middle Ages	Futures II
2	Fabio Zampieri	The Holistic Approach of Evolutionary Medicine: An Epistemological Analysis	Futures II
3	Lynnette Leidy Sievert	Choosing the Gold Standard: Subjective Report vs Physiological Measure	Futures II
4	Elizabeth Edwards	Photography, Survey and the Desire for 'History'	Futures II
5	Ben Anderson	Emergency Futures	Futures
6	Pier Paolo Saviotti	Are There Discontinuities in Economic Development?	Futures II
7	Sander L. Gilman	'Stand Up Straight': Notes Toward a History of Posture	Futures II
8	Meredith Lloyd-Evans	Limitations and Liberations	Futures II
2013 Volume 6			
1	David Martin-Jones	The Cinematic Temporalities of Modernity: Deleuze, Quijano and *How Tasty was my Little Frenchman*	Time
2	Robert Levine	Time Use, Happiness and Implications for Social Policy: A Report to the United Nations	Time
3	Andy Wood	Popular Senses of Time and Place in Tudor and Stuart England	Time
4	Robert Hannah	From Here to the Hereafter: 'Genesis' and 'Apogenesis' in Ancient Philosophy and Architecture	Time
5	Alia Al-Saji	Too Late: Racialized Time and the Closure of the Past	Time
6	Simon Prosser	Is there a 'Specious Present'?	Time
2014 Volume 7			
1	Robert Fosbury	Colours from Earth	Light
2	Mary Manjikian	Thinking about Crisis, Thinking about Emergency	Time
3	Tim Edensor	The Potentialities of Light Festivals	Light
4	Angharad Closs Stephens	National and Urban Ways of Seeing	Light
5	Robert de Mello Koch	From Field Theory to Spacetime Using Permutations	Time
6	Jonathan Ben-Dov	What's In a Year? An Incomplete Study on the Notion of Completeness	Time
7	Lesley Chamberlain	Clarifying the Enlightenment	Light
8	Fokko Jan Dijksterhuis	Matters of Light. Ways of Knowing in Enlightened Optics	Light
2015 Volume 8			
1	Valerie M. Jones	Mobile Health Systems and Emergence	Emergence
2	Stéphanie Portet	Studying the Cytoskeleton: Case of Intermediate Filaments	Modelling
3	Peter Cane	Two Conceptions of Constitutional Rights	Emergence
4	Nathan J. Citino	Cultural Encounter as 'Emergence': Rethinking US-Arab Relations	Emergence
5	N. Katherine Hayles	Nonconscious Cognition and Jess Stoner's *I Have Blinded Myself Writing This*	Emergence
6	Alice Hills	Waiting for Tipping Points	Emergence
7	Margaret Morrison	Mathematical Explanation and Complex Systems	Emergence
8	Tim Thornton	Emergence, Meaning and Rationality	Emergence
9	John Heil	The Mystery of the Mystery of Consciousness	Emergence

No.	Author	Title	Series
10	David C. Geary	Sex Differences in Vulnerability	Emergence
11	Richard Read	Negation, Possibilisation, Emergence and the Reversed Painting	Emergence

2016 Volume 9

1	George Williams	An Australian Perspective on the UK Human Rights Act Debate	Evidence
2	James E. Gardner	Can We Gain Evidence About Volcanic Pyroclastic Flows from Those Who Survive Them?	Evidence
3	John Brewer	Art and the Evidence of Attribution. Giovanni Morelli, Morellians and Morellianism: Thoughts on 'Scientific Connoisseurship'	Evidence
4	Claire Langhamer	An Archive of Feeling? Mass Observation and the Mid-Century Moment	Evidence
5	Heike Egner	The IPCC's Interdisciplinary Dilemma: What Natural and Social Sciences Could (and Should) Learn from Physics	Evidence
6	Barbara Dancygier	Reading Images, Reading Words: Visual and Textual Conceptualization of Barriers and Containers	Evidence

Insights

Insights is edited by Nicholas Saul, IAS Director and Professor of German Literature and Intellectual History.

Correspondence should be directed to Pauline Edmondson (pauline.edmondson@durham.ac.uk).