Chapter 6
Relevance Realization and the Neurodynamics and Neuroconnectivity of General Intelligence

John Vervaeke and Leonardo Ferraro

Abstract In this paper we review arguments for the central nature of the problem of relevance, as well as arguing that relevance realization is the basis for general intelligence, supporting this position with recent findings in neurodynamics and neuroanatomy, as well as machine learning and graph theory.

6.1 Introduction

This paper will present five linked arguments. The first argument will outline the centrality of the problem of relevance, and how it is becoming the focus of an emerging framework in cognitive science. The second argument will explore some of the central features needed in an account of relevance. This exploration will lead to our third argument, namely that there cannot be a scientific theory of relevance. However, this will not be cause for despair, because once we abandon a search for an account of relevance we can successfully pursue a theory of relevance *realization*. The fourth argument will outline such a theory in terms of the *bioeconomics* of relevance realization. Finally, the fifth argument will show how this theory of bioeconomical relevance realization provides a basis for an explanation of general intelligence in terms of the neurodynamics and neuroconnectivity in the brain. This set of arguments will help to reveal the self-organizing and plastic nature of general intelligence in a way that would lay the foundations for the autonomous agents that are central to the SmartData vision.

J. Vervaeke (✉)
Cognitive Science, University of Toronto, Toronto, ON, Canada
e-mail: john.vervaeke@gmail.com

L. Ferraro
University of Toronto, Toronto, ON, Canada
e-mail: leonardo.ferraro@utoronto.ca

I. Harvey et al. (eds.), *SmartData: Privacy Meets Evolutionary Robotics*,
DOI 10.1007/978-1-4614-6409-9_6,

6.2 The Centrality of the Problem of Relevance

Vervaeke et al. [19] argued that the problem of how agents zero in on relevant information was emerging as the central issue driving many different difficulties in cognitive science. The paper reviewed dilemmas within problem solving, categorization, communication, robotic interaction, and rationality to show that all these issues converged on the problem of how cognitive agents determine the relevance of information. Here, we will briefly review the case for the centrality of the problem of relevance by exploring the literature on problem solving, an ability that is central to being an autonomous intelligent agent.

Three related areas within the psychology of problem solving articulate the central importance of relevance. These are the issues of combinatorial explosion, the ill-definedness of real-world problems, and the need for insight within problem solving. Combinatorial explosion was revealed in the seminal formalization of problem solving by Newell and Simon [13] in which a problem is represented by an initial non-desirable state, operations that can transform that state into other states, a desired goal state that should be the end result of the transformations, and path constraints on how one was allowed to perform sequences of operations. So, for example, an initial state could be one of hunger, with the operations being things one can do to alter one's state, such as walking or throwing something. The goal state would be the ending of hunger while path constraints might include that one is not allowed to end hunger by killing oneself, or burning down one's house in order to cook all the food in it. A problem solution is a sequence of operations that transforms an initial state into the goal state while obeying the path constraints. The set of alternative possible pathways of transformations can be represented by a search space.

It is important to remember that in real life one does not have the god's eye point of view that reveals which pathway of operations is correct. Additionally, one cannot search the whole space to determine the correct pathway because the number of alternatives available is extremely vast. The formula for calculating the number of pathways is F^D, where F represents the number of operations available to one, and D is the number of steps one takes. So, for example, in a typical game of chess one can make 30 legal moves and one takes about 60 turns. So the number of alternative pathways to checkmate is 30 to the power of 60 or 4.239×10^{88}. Compare that to the number of neurons in the brain (estimated to be 10^{10}), or even the number of synaptic connections (approximately 5×10^{14}). In fact, the number of atoms in the universe is 10^{82}. So even the massively parallel nature of the brain is not sufficient for searching the entire search space using a brute force, exhaustive strategy; the size of the search space is just too vast. There is a *combinatorial explosion* in the number of alternatives one has to check, one that requires us to somehow home in on worthwhile paths and ignore others.

As Cherniak [3] famously noted, we are in the finitary predicament, in that we have limited time and resources with which to solve our problems. Our search must be a *heuristic* one that does not exhaust all of the available alternatives.

Instead, a heuristic biases the search to a restricted area within the total search space. The use of heuristics such as means-ends analysis was the solution to combinatorial explosion proposed by Newell and Simon. The problem is that while heuristics are necessary for addressing combinatorial explosion, they are not sufficient for doing so. They are insufficient because they require a pre-specification of what area of the search space to check, and yet it often happens that this pre-specification does not match up to the problem at hand. The success of the pre-specification also depends upon the size of the search space. If the search space is very large, as it often is, then even heuristic search can be very time consuming. The problem is that these two concerns are in a trade-off relationship. As we open up the pre-specification within the heuristic so it is more likely to apply to the problem at hand and thereby succeed, we also dramatically increase the amount of time needed to apply it.

Somehow problem solvers reliably (but not perfectly) zero in on the relevant information to be investigated. They do not do this abstractly, but in the way they formulate the problem. How individuals represent the initial state, goal state, operations, and path constraints, is the way in which they attempt to zero in on relevant information. Problem formulation is how problem solvers constrain the search space of a particular problem so that heuristics can effectively apply to it. Problem formulation captures what problem solvers deem relevant to a specific problem, and this formulation helps them to intelligently ignore most of the information in the search space. To intelligently ignore means that problem solvers do not even consider most available information, and they find and focus upon information that turns out to be relevant without comparing it to all the irrelevant information that is available.

However, problem formulation also addresses another core difficulty facing problem solvers, viz., most real world problems are ill-defined problems. Unlike chess, where the initial state, goal state, operations, and path constraints are clear and helpfully represented, ill-defined problems lack such clear representations. With ill-defined problems, the goal state is often murky, the initial state is unclear, and the operations are unspecified. So, for example, writing a good paper is an ill-defined problem: the goal state is unclear, in that the properties of a good paper do not seem to be readily accessible. Note that one is tempted to answer this by using synonyms for relevance such as "a good paper presents *important* information, a good paper covers *key* material in a *succinct* manner," etc. The initial state is not having a good paper, but what should one pay attention to in this state in order to provide guidance? Saying that one should pay attention to similar solutions in the past is not helpful for two reasons. Firstly, similarity presupposes relevance realization abilities, in that, trivially-speaking, all things are infinitely similar and dissimilar [6]: we can show that two seemingly distinct objects (say, your kitchen table and your car) are both smaller than Jupiter, both located in the same continent, etc. As such, similarity cannot be based simply upon shared features, but rather they must be based on a set of shared *relevant* features. Also, the previous solutions are similar precisely in terms of their ability to deal with ill-definedness, and so we face the problem of how to get successful problem solving started in the first place.

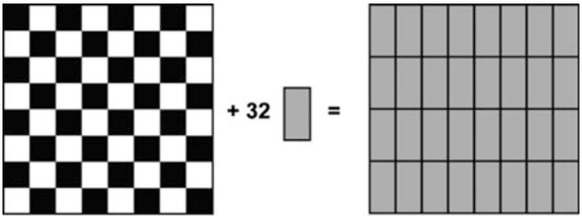

Fig. 6.1 Setting up the mutilated chessboard problem

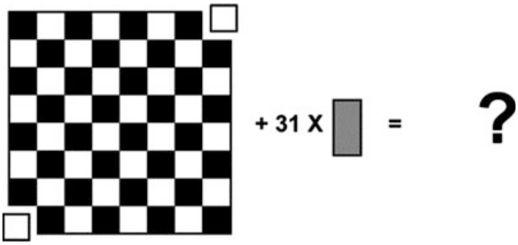

Fig. 6.2 Can the mutilated chessboard be covered?

Consider how we are beset by ill-defined problems all day long such as following and/or joining a conversation successfully, getting and/or telling a joke or telling a story. As such, we frequently have to generate problem formulations in order to both address ill-definedness and to avoid combinatorial explosion. Not only do we have to zero in on relevant information, we often have to generate missing relevant information in order to do so. Problem formulation handles both of these demands and is a primary way in which we do relevance realization within problem solving.

However, sometimes the problem formulation constrains a problem in such a way so that it cannot be solved; the problem formulation itself becomes problematic. In these circumstances, the problem formulations need to be broken up and reformulated. The solver must have an insight: their relevance realization abilities have to be flexibly recursive and self-correcting. They have to realize that they have misformulated the problem and be capable of generating a new relevance profile that then informs a new, more effective formulation. This is made evident in the work done by Kaplan and Simon [9] on the mutilated chessboard problem (see Fig. 6.1).

Consider a standard chessboard of 64 squares. Consider also a domino that will cover two squares either vertically or horizontally. Can one cover the board, without overhang or overlap, with 32 dominos? Most people answer this easily with a "yes." However, now consider if the chessboard is mutilated so that two diagonally opposite corner squares are removed (see Fig. 6.2). Can the remaining 62 squares be covered by 31 dominos without overhang or overlap, with a proof that the answer is correct? Most people formulate this as a covering problem, trying to visualize different patterns of covering the mutilated board with dominos.

They are quickly overwhelmed by combinatorial explosion, and it is unclear if they are bringing the relevant operations to bear upon the problem. However, if one notices, i.e. finds salient and makes relevant, the fact that the two corner pieces are always the same colour, then one may have an insight. Each domino always covers a black square and a white square, no matter its orientation. Therefore, an equal number of black and white squares are needed in order for the dominos to cover the board without overlap. Since the two corner squares are the same colour, the even black to white square ratio is broken and thus it is necessarily the case that the board cannot be covered with the 31 dominos.

As we can see, relevance realization is crucially important to our problem-solving abilities. Similarly, it is essential in many other areas, such as categorization and communication (see [19] for a more detailed treatment of each). More than being significant within each of these central abilities of intelligence, it is also found between them. This is because the abilities are inter-dependent: problem solving requires good categorization and communication (at least within the individual across time and memory), and communication requires good categorization and problem solving abilities, etc. The relevance realization that we do within each crucial process must be relevant to the relevance realization going on in the other central processes. The ability to determine relevance is foundational to our intelligence and would be central to any SmartData system.

6.3 The Necessary Features of an Account of Relevance

The question now arises as to where within the mind/brain should we look for the processing of relevance. What is the correct level of analysis? Perhaps relevance is a property of representations the way truth is such a property. However, this cannot be correct [19] because, as Searle [15] has pointed out, all representation is aspectual. We never represent all the properties of a thing because we are in the finitary predicament. We always only represent a subset of any real world things. This subset of features and how they hang together is an aspect. All representation is aspectual. Yet an aspect is a zeroing in on properties deemed relevant, and a formulation of those properties as highly relevant to each other and to oneself and others. The ability to represent crucially presupposes the ability to realize relevance and therefore cannot serve as a basis of explaining it without circularity. This has the very important consequence that relevance realization cannot ultimately rely upon or begin with the brain representing certain external states or goals in the world. Relevance realization has to initially be completely internal to the brain.

Perhaps the computational level within the brain's information processing is the correct level of analysis. The computational level is the level at which information is encoded in logically structured propositions and manipulated in a rule-governed, inferential manner. However, there are two important arguments against situating relevance realization at this level. The first comes from Fodor [4], which is surprising since he is one of the staunchest defenders of a computational theory of mind. The problem, according to Fodor, is that relevance is an issue of cognitive commitment. It is how much of your limited attention, time and, resources you are going to give to something. This cognitive commitment depends on the current context and one's idiosyncratic history of previous commitments. That cannot be captured in the syntactic/logical structure of a proposition because we need that to be invariant across situations and people in order to function within truth preservation and generation. We cannot pre-specify the commitment *to* the proposition *within* the structure of the proposition. Yet, all the rules governing the inferential manipulation of the proposition work solely in terms of the invariant constitutive structure. We can find a proposition relevant one moment and completely irrelevant the next, even though the proposition, its structure, and the rule that govern it have not changed. Wittgenstein et al. [20] argued that every rule, in order to be used, has to be interpreted and specified. This interpretation and specification cannot itself be captured in rules on pain of infinite regress. The process must bottom out in a process that is not itself rule governed, and these processes of interpretation and specification are processes laden with relevance realization. If we put Fodor and Wittgenstein's arguments together we get that the computational level does not capture cognitive commitment and presupposes relevance realization in its use of rules. For these reasons it is not the correct level at which to explain relevance realization. Relevance realization must be happening at a level of analysis more basic than standard information processing. Yet it cannot be handled by some central processor because that processor would face a combinatorial explosion of information facing it even within the brain's own processing. The processing of relevance realization has to happen as a constraint on all processing both local and global within the brain. Relevance realization has to be internal, sub-semantic, sub-syntactic, and scale invariant in its operations. Finally, it must be completely self-organizing because it has to be a self-correcting and self-transforming process.

Vervaeke et al. [19] argue that *economical* properties best satisfy these requirements of relevance realization. These are *logistical* properties that concern decisions about how to commit resources and ration time and processing, rather than *logical* properties governing truth preservation within inference. Important logistical properties are efficiency, which operates on metabolic expenditure and the obtaining of reward, and optimization functions on the attaining of said rewards. These logistical properties are internal to the biology of the organism. Hence these properties should more properly be called bioeconomical properties. Bioeconomical properties are self-organizing and scale invariant. They are sub-semantic and sub-syntactic, completely internal and vital to the biology of the organism.

Bioeconomical processing results in a brain that dynamically couples to its environment in a way that results in intelligent behaviour. For example, intelligent behaviour requires a dynamical equilibrium between exploiting current sources of reward and exploring for better opportunities. One way of improving how a system obtains reward is if it gives a weighting on a behaviour's prediction of reward based on temporal lag between the behavioiur and the reward. This is called *temporal displacement learning* (see [19] for more discussion). In contrast, a system can improve its chance on reward if it has *inhibition on return*, which causes a system to avoid repeated use of the same stimulus (see [19] for more discussion). If a brain internally pits temporal displacement learning, which reinforces behaviour thereby driving further exploitation, and inhibition on return, which will drive exploration, then the organism with such a brain will flexibly exploit and explore its environment without specifically setting or possessing goals of either exploitation or exploration. This results in a continual self-organization of behaviour and development that precludes any homunculi or chicken and egg problems. No separate central executive is required to make these decisions. The decisions emerge out of the self-organized processing of the bioeconomical properties (see [14] on emergent activity switching).

6.4 Relevance vs. Relevance Realization

However, there is now a difficulty. One may try to create a theory of relevance in terms of such bioeconomical properties, but there can be no scientific theory of relevance. In order to generate the inductive generalizations that are central to science, scientific reasoning requires classes that support such induction. Such classes require that its members possess homogenous, stable and intrinsic properties. We cannot have a science of things that happen on Tuesdays because the set of events is not homogenous, stable, nor are Tuesdays intrinsic to the world. Similarly, the things we find relevant do not form homogenous, stable classes, nor is relevance intrinsic to the world. So we cannot have a scientific theory of relevance.

Yet this is not cause for despair. Consider an analogy: we cannot form a scientific theory of Darwinian fitness because the set of features that makes a creature fit is not homogenous, nor stable, nor intrinsic to the biology of the creature. What Darwin's theory gave us was an account of how fitness was continually being redesigned in a self-organizing and contextually sensitive manner. So we do not need a theory of relevance; we only need a theory of *relevance realization*. We need a theory of how cognition continually redesigns itself to fit the changing world.

What Darwin also gave us was a mechanism for evolution. He proposed a virtual governor [8] in which there is a configuration of enabling and selective constraints. Enabling constraints, such as mutation and sexual reproduction, generate options, while selective constraints, such as competition and environmental disaster, winnow them down. Vervaeke et al. [19] proposed that the mechanism of relevance realization was just such a virtual governor, operating on the bioeconomical properties of cognition. There are constraints of efficiency that put selective pressure on processing while there are constraints of resiliency that enable new possibilities of processing. Cognition evolves its fitness to its environment in a dynamical self-organizing manner.

Remember that no heuristic operates well across all domains because it attempts to pre-specify where to search for relevance. So the price paid for domains in which it enhances performance is the detriment to processing in other areas. However, evolution has created a solution to this problem: it finds heuristics that are in a trade-off relationship with each other and then puts them into opponent processing, functionally integrating them into a push/pull relationship as they pursue their opposed goals. So, for example, the parasympathetic and sympathetic nervous system use such opponent processing to continually redesign the level of arousal in a contextually sensitive manner. The brain is pursuing two logistical properties that are in such a trade-off relationship, viz., efficiency and resiliency. The notion that brains are processing information in a way that is governed by efficiency is a view that is becoming central in cognitive science (see [5, 12, 17]). Less explored is the idea that brains are also seeking the opponent goal of resiliency. Brains are trying to maintain an important degree of flexibility so that they have the potential to redesign their function, thereby increasing their fault tolerance in order to retain a potential to resist damage. Thus, the brain can have opponent processing between efficiency and resiliency function as a virtual governor that sets parameters on cost functions that optimize for reward. This would be the machinery of relevance realization.

6.5 From Relevance Realization to Cognitive Development and General Intelligence

Vervaeke et al. [19] proposed that this virtual governor consists of nested virtual governors that carry out more specific opponent processing between efficiency and resiliency (see Table 6.1).

So, for example, one such nested governor (see [19] for more) is one that performs opponent processing between data compression for efficiency and data particularization for resiliency. In data compression one is doing something analogous to finding the line of best fit for data, while in particularization one is allowing the function to move towards over-fitting to the data. A brain that is doing this internally will dynamically couple to its world in a way that is always trading off between being an efficient general purpose machine and being a resilient set of special purpose machines. Neither strategy is comprehensively fit, but to continually shift between them is. The brain is not trying to be either type of machine; the type of machine it becomes results from the coupling of its internal processing to both cross contextually invariant patterns, tracked by compression, and more contextually specific patterns which are tracked by particularization. In this way it manages the *applicability* of its information. We have discussed how opponent

Table 6.1 Mapping bioeconomics onto behavioural repertoire

Internal bioeconomic property	External interactive property
Cognitive scope (compression vs. particularization)	Applicability (general purpose vs. special function)
Cognitive tempering (temporal displacement learning vs. inhibition of return)	Projectability (exploiting vs. exploring)
Cognitive prioritization (cost function #1 vs. cost function #2)	Flexible gambling (focusing vs. diversifying)

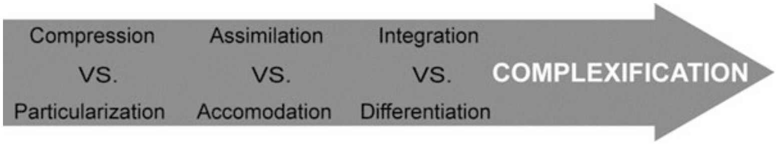

Fig. 6.3 The emergence of complexification from opponent processing

processing between temporal displacement learning and inhibition on return affords the management of exploitation, which is efficient and exploration, which introduces resilience. Vervaeke et al. [19] called this the *projectability* of information. In addition, the brain must trade between different channels of rewards and the cost functions that try to optimize the obtaining of reward. The brain thereby flexibly decides if it should gamble by focusing all its efforts on one or a few channels in the hopes of a big payoff, or diversifying its efforts to hedge its bets. In this way, the brain prioritizes its cost functions in a self-organizing manner that again trades between efficient expenditure in focusing and expenditure that introduces resiliency through diversification.

For the sake of furthering the primary argument connecting relevance realization and general intelligence, we will now focus primarily on compression vs. particularization (nevertheless, we predict that the other governors will also be found to be predictive of general intelligence as measured by psychometric tests). Compression results in the assimilation of information to existing structures and therefore results in the integration of information. Particularization results in the accommodation of existing structures to information so that the differentiation of information occurs. Since the brain is doing both in an opponent fashion, it is simultaneously developing both integration and differentiation. A system that simultaneously integrates and differentiates its functions is complexifying as a system (see Fig. 6.3).

This is important because complex systems have emergent functions. This is the way in which the brain can develop its competence to deal with a complex world, viz., it self-complexifies (or *develops*).

We, in conjunction with Zachery Irving (Irving, Z., Vervaeke, J., & Ferraro, L. (2010) The Relevance Realization Framework of Intelligence: Integrated evidence from cognitive science, psychometrics, and neurodynamics. Unpublished manuscript), have argued that relevance realization is central to those abilities that make one a cognitive agent, and that those abilities are also those that are measured by psychometric tests (see Fig. 6.4).

It is well established that such psychometric tests show a positive manifold in which performance on each test is strongly predictive of performance on the other tests [7, 16]. This strongly suggests a central underlying ability often called general intelligence. It is therefore extremely plausible that general intelligence is the central underlying ability of relevance realization (see Fig. 6.5).

This, in turn, suggests that processes in the brain that can plausibly be interpreted as performing the opponent processing between compression and particularization should be predictive of general intelligence. A scale invariant [1, 10],

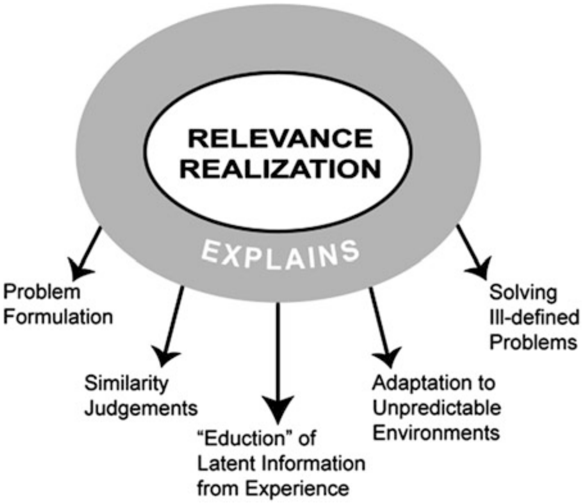

Fig. 6.4 The explanatory scope of relevance realization

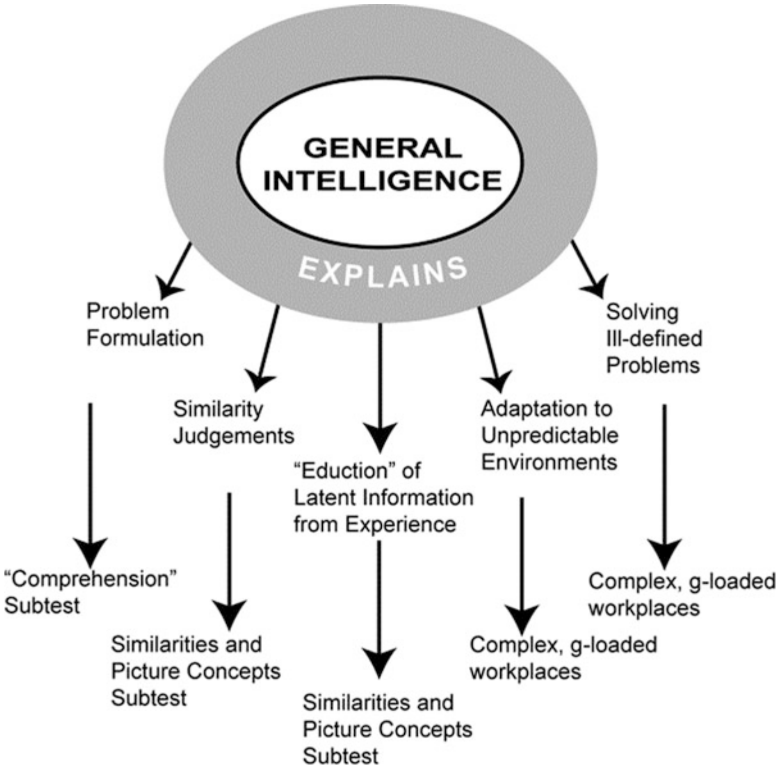

Fig. 6.5 The explanatory congruence of relevance realization and general intelligence

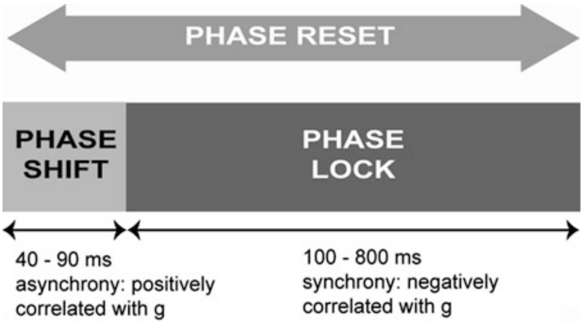

Fig. 6.6 The components of phase reset

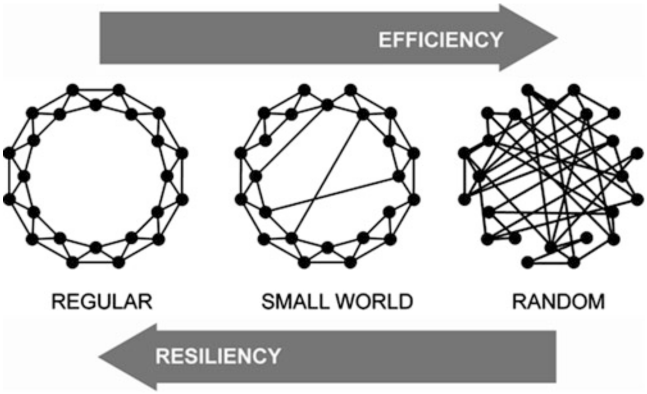

Fig. 6.7 Contrasting network topologies

self-organizing process of patterns of neuronal firing in the brain has been discovered. The brain's neuronal firing goes through self-organizing criticality (SOC) in which it oscillates between patterns of synchronous firing of neurons and periods of asynchronous firings. The synchronous firing probably is carrying out information integration by compression while the asynchronous period affords the differentiation and specialized processing of separate groupings of neurons. If this is correct we can predict that variations in the flexibility of SOC in brains should correlate with variations in measures of general intelligence. Thatcher et al. [18] have found exactly this (see Fig. 6.6).

The neurodynamics of brain firing seem to instantiate the machinery of relevance realization in order to afford general intelligence.

Not only the brain's firing but its wiring should also show evidence of being governed by relevance realization machinery. The brain has been shown to wire into small world networks in a scale invariant manner [2]. Small world networks show features of both regular networks that are highly resilient, and random networks that are highly efficient (see Fig. 6.7).

Langer et al. [11] have recently shown that the more a brain wires in a small world network fashion, the more intelligent it is. In both its firing and its wiring, the brain is pursuing a trade-off between efficiency and resiliency, and this results in the brain possessing general intelligence. If SmartData is about autonomously intelligent agents then it is about virtual agents that will possess general intelligence. They will do this by instantiating the virtual governors of relevance realization. These governors in turn can be implemented in a neurodynamics of self-organizing criticality and neuroconnectivity of small world networks. By creating virtual versions of this firing and wiring of the brain, virtual agents can realize relevance in an on-going and evolving manner, and thereby become truly SmartData.

Acknowledgement We would like to thank Zachery Irving for his collaboration on a previous manuscript on relevance realization and general intelligence. His work on that paper help to generate ideas for this paper.

References

1. Bak P, Tang C, & Wiesenfeld K (1988). Self-organized criticality. *Phys Rev A*, *38*(1), 364–374.
2. Bullmore, E., & Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat Rev Neurosci*, *10*(3), 186–198.
3. Cherniak, C. (1990). *Minimal Rationality*. MIT Press.
4. Fodor, J. (2006). How the mind works: what we still don't know. *Daedalus*, *135*(3), 86–94.
5. Friston, K. (2010). The free-energy principle: a unified brain theory? *Nat Rev Neurosci*, *11*(2), 127–138.
6. Goodman, N. (1972). *Problems and projects*. Bobbs-Merrill.
7. Jensen, A. R. (1998). *The g factor: the science of mental ability*. Praeger.
8. Juarrero, A. (2002). *Dynamics in Action: Intentional Behavior As a Complex System*. MIT Press.
9. Kaplan, C. A., & Simon, H. A. (1990). In search of insight. *Cognitive Psychol*, *22*(3), 374–419.
10. Kitzbichler, M. G., Smith, M. L., Christensen, S. R., & Bullmore, E. (2009). Broadband Criticality of Human Brain Network Synchronization. *PLoS Comput Biol*, *5*(3), e1000314.
11. Langer, N., Pedroni, A., Gianotti, L. R. R., Hänggi, J., Knoch, D., & Jäncke, L. (2012). Functional brain network efficiency predicts intelligence. *Hum Brain Mapp*, *33*(6), 1393–1406.
12. Montague, R. (2007). *Your Brain Is (Almost) Perfect: How We Make Decisions*. Plume.
13. Newell, A., & Simon, H. A. (1972). *Human problem solving*. Prentice-Hall.
14. Perkins, D. N. (2002). The engine of folly. *Why smart people can be so stupid* (pp. 64–85). New Haven, CT, US: Yale University Press.
15. Searle, J. R. (1992). *The Rediscovery of the Mind*. MIT Press.
16. Spearman, C. (1904). "General Intelligence," Objectively Determined and Measured. *The Am J Psychol*, *15*(2), 201–292.
17. Sperber, D., & Wilson, D. (1995). *Relevance: Communication and Cognition*. Blackwell.
18. Thatcher, R. W., North, D. M., & Biver, C. J. (2009). Self-organized criticality and the development of EEG phase reset. *Hum Brain Mapp*, *30*(2), 553–574.
19. Vervaeke, J., Lillicrap, T. P., & Richards, B. A. (2012). Relevance Realization and the Emerging Framework in Cognitive Science. *J Logic Comput*, *22*(1), 79–99.
20. Wittgenstein, L., Hacker, P. M. S. (ed), & Schulte, J. (ed) (1953/2009). *Philosophical Investigations*. John Wiley & Sons.