Draft: to be published as:
Heylighen, F. (2019). Transcending the Rational Symbol System: how information technology integrates science, art, philosophy and spirituality into a global brain. In N. Gontier (Ed.), *Handbook of Human Symbolic Evolution*. Oxford University Press.

Transcending the Rational Symbol System: how information and communication technology integrates science, art, philosophy and spirituality into a global brain

Francis Heylighen
Center Leo Apostel,
Vrije Universiteit Brussel
fheyligh@vub.ac.be

Abstract:
Because symbols support the uniquely human capabilities of language, culture and thinking, cognitive science has tried to explain intelligence as founded on rational symbol systems (RSS). Such systems consist of collections of symbols along with the logical and grammatical rules required for combining these symbols into meaningful expressions. The main shortcoming of this RSS mechanism is that it reduces the continuous experience of reality to a combination of static, discrete, and to some degree arbitrary, elements. To fully support intelligence, such symbols need to be grounded in subsymbolic networks of associations and situated interactions that capture the subjective, transient and intuitive aspects of experience. Historically, different approaches have tried to overcome the shortcomings of rational symbol systems, albeit as yet with limited success. These include science, by formalizing and operationalizing symbols; philosophy, by seeking for the reality behind symbols while analyzing the shortcomings of symbolic representations; art, by evoking intuitive insights and experiences; and spirituality, by expanding consciousness beyond rational symbol systems. The on-going explosion in information and communication technology now makes it possible to extend and integrate their results, through techniques such as computer simulation, artificial intelligence, neural networks, multimedia, virtual and augmented reality, and brain-computer interfaces. Thus, the subsymbolic level of cognition, like the symbolic level before it, becomes externalized and controllable, supporting and automating creativity and intuition. It is proposed that this will produce an evolutionary transition to a supra-human level of knowledge, intelligence and consciousness, envisioned as a Global Brain for humanity.

Introduction

A *symbol* is a token, such as a word, picture or gesture, that conventionally stands for or represents something else: the symbol's meaning, reference or denotation (Burks, 1949; Nauta, 1972; Chandler, 2017). The use of symbols is arguably what made humans different from the animals out of which they evolved. Symbols gave us the language we use to not only communicate, but reason. Symbols enabled us to accumulate and propagate knowledge and culture—not just across populations, but across generations. The resulting growth of knowledge culminated in the extremely sophisticated science, technology and culture of the 21st century, which made humans into the dominant species on the planet.

Symbols made this possible by isolating specific items in our mind, such as concepts, thoughts or feelings, and externalizing them in the form of observable patterns. Such patterns represent the mental content while being imprinted onto a manipulatable, physical carrier. For speech, the carrier is sound; for pictures, it is a screen or canvas; for text it is paper or silicon memory circuits. This externalization helped us to preserve the otherwise fleeting content of our consciousness. Thus, we could register an accurate and enduring memory of potentially important information (Fig. 1). It also allowed us to transmit that information to others, so that they too could benefit from it—and potentially improve on it by adding their own information. It finally enabled us to manipulate or process that information, by combining different symbols into new expressions, thus producing as yet unseen content.

This was probably the essential step in the emergence of our typically human form of intelligence. What really makes human cognition so powerful compared to animal cognition is our ability to conceive of situations that we have never experienced in reality. That is because we can represent such as yet unknown situations by novel combinations of known symbols. In the simplest case, we receive such a representation as a communication from someone else, across space or time. For example, we can get an idea of how a tool we never saw functions, by listening to someone's explanation, reading a description, or studying a diagram. But symbols also enable creative thinking: they allow us to conceive of a tool that does not exist yet, to reflect on what we would do in a hypothetical situation, or to design a building that will take years to construct.

How symbols enable endless creativity was perhaps explained most clearly by the linguist Chomsky's concept of *generative grammar* (Chomsky, 2014; Horrocks, 2014). This is a system of rules that allows you to produce grammatically correct (i.e. understandable) expressions by combining words (i.e. symbols) from a lexicon. The fundamental insight is that both the lexicon and the collection of rules are finite, but the set of expressions that can be generated from them is infinite. Thus, language allows us to recursively produce an infinite number of sentences and sequences of sentences, potentially describing an infinite number of situations. (Note that Chomsky assumed that the human brain possesses an innate "language organ" lacking in animals that enables such recursive combinations. However, this hypothesis is no longer considered very plausible, and is anyway not needed to explain how symbols enable creativity). The generative rules of grammar are complemented by the constraining rules of *logic*, which tell us which of the generated expressions represent possible situations, and how we can infer from them further propositions about these situations.

Fig. 1: symbols, such as the words in a book, allow us to externalize mental content, and thus dependably store, communicate and manipulate it.

This mechanism of generating abstract representations and then reasoning about them—even in the absence of the phenomena being represented—allows us to reflect, plan, and solve problems. Thus, it appears like the essence of intelligence. Therefore, Newell and Simon, two of the founders of cognitive science and artificial intelligence, concluded that a collection of symbols together with a collection of rules for manipulating them is all you need to produce intelligent behavior. They formulated this assumption about cognition as the “physical symbol system hypothesis” (Nilsson 2007; Newell 1980; Newell & Simon 1976).

This idea supported a further step in the externalization of mental content: thanks to computers we no longer need people to manipulate symbols. If we program a computer to apply the appropriate rules, then this program can do the work for us, and reason intelligently about whatever situation can be described with symbols. This generalized the notion of *information processing:* represent the information you need to reason about as a combination of symbols (words, characters, numbers, or ultimately 1s and 0s), enter it into a suitably programmed machine, and then you can delegate the desired further manipulation of that information to the machine. This provided the foundation for what we now call *Information and Communication Technology* (ICT). The main components of ICT are memories to store the combinations of symbols, communication links to transmit these combinations to other memories, and processors (hardware) and programs (software) to transform the combinations into new, potentially more informative combinations.

This externalization and automation of memory, communication, and processing spectacularly enhanced the emergence of *distributed intelligence* or *distributed cognition* (Dror & S. R. Harnad 2008; Fischer 2006; Hutchins 2000). Previously, processes like memorizing, remembering, and thinking largely happened inside an individual person's brain, supported at most by external documents or verbal communications. ICT made it possible to efficiently distribute these processes across a large array of communicating components, including people, computers, software, robots, sensors, and remotely controlled machines. Each of those performs that part of the process for which it is best suited, while the communication links between them allow them to coordinate their activities. Nowadays, the Internet allows such components to collaborate so effectively that they become capable of functioning like a single cognitive system at the planetary scale—a "Global Brain" (B. Goertzel 2002; Heylighen 2011; Mayer-Kress & Barczys 1995). This leads us to envisage a distributed intelligence that surpasses human intelligence as radically as human intelligence surpasses animal intelligence. The present article will explore the benefits and limitations of different types of symbol systems in supporting such a superhuman intelligence.

We will in particular examine the shortcomings of the traditional, "rational" or "verbal" form of symbolization, which reduces an infinitely complex, fluid and changing reality to a finite collection of discrete and static symbols ("words") and rules ("grammar" and "logic") (Shands 1971). We will call such a collection a *Rational Symbol System* (RSS). (Note that this replaces my earlier term "conceptual-symbolic code" (Heylighen 1984)). The RSS term seems to better capture the abstract, rule-based reasoning mode of this form of representation than Newell and Simon's (1976) term "physical symbol system".

Science, philosophy, art, and spirituality all have to some degree attempted to overcome the limitations of rational symbol systems by expressing meanings for which no adequate symbols or language exist (Heylighen, 1984). ICT is now lending them a helping hand, by proposing forms of representation, such as neural networks (McLeod, Plunkett and Rolls, 1998; Schmidhuber, 2015), multi-agent simulations (Kinny and Georgeff, 1997; Ferber, 1998), virtual reality (Krueger, Packer and Jordan, 2001) and augmented reality (Van Krevelen and Poelman, 2010), that are much richer, more dynamic and adaptive than verbal or logical descriptions. We will here not go into the technical details that characterize these novel tools, but focus on how they may help us to transcend the limitations of human intelligence, and thus prepare the grounds for a supra-human form of cognition. In that way, we may extrapolate human symbolic evolution to the next "major evolutionary transition" (Maynard Smith and Szathmáry, 1997) or "metasystem transition" (Turchin, 1977, 1995; Heylighen, 1995) that would lead to a higher level of intelligent organization.

Symbolic Cognition and its shortcomings

The study of intelligence as a problem-solving mechanism began in the 1950's under the influence of the newly emerging approaches of cybernetics (Ashby, 1964; Wiener, 1965; Heylighen and Joslyn, 2003), systems theory (Boulding, 1956; Mesarović and Takahara, 1975) and information theory

(Shannon and Weaver, 1963; Nauta, 1972). These conceived of an intelligent agent as a system that receives an input of information that represents some problem to be solved or question to be answered. The agent then internally processes that information, using specific routines, rules or programs, in order to transform it into an output that represents a solution to the problem or an answer to the question. This perspective was applied to understand human perception, problem solving (Newell and Simon, 1972) and decision making as the processing of information. This approach initiated the domain of cognitive psychology, and more generally cognitive science (Thagard, 1996; Bechtel and Graham, 1999).

The same perspective inspired researchers to program computers, thus initiating the domain of Artificial Intelligence (AI) (Goertzel and Pennachin, 2007; Russell and Norvig, 2009), which tries to simulate human intelligence in software. Computers had hitherto been used just to make calculations, i.e. to process numbers. Tackling general problems required not just quantitative, but qualitative descriptions—using symbols equivalent to words, or what in logic are called 'subjects' and 'predicates'. The rules of grammar and logic seemed perfectly suited to process such symbolic expressions, by specifying how a particular expression (i.e. combination of symbols) should be transformed into another combination.

This understanding of cognition as the rule-governed manipulation of discrete symbols was initially very promising. Rational symbol systems (RSS) not only seemed able to explain the production and understanding of language, logical reasoning, and problem solving in humans; they also allowed us to simulate these processes with computers (Newell and Simon, 1976). AI researchers quickly produced a number of programs that could perform apparently intelligent activities, like playing chess, proving theorems, or diagnosing diseases from lists of symptoms. However, after these early successes, by the late 1980's progress in AI seemed to stall. It turned out that the computer simulation of human intelligence in more concrete, real-world domains was much more difficult than expected. For example, symbolic AI programs were hopelessly unreliable in performing apparently simple tasks like translating between natural languages, recognizing objects in visual scenes, or directing the actions of an autonomous robot. This resulted in increasingly stringent critiques of the symbolic paradigm for understanding cognition (Dreyfus, 1992; Bickhard and Terveen, 1995).

Perhaps the most fundamental shortcoming pointed out was the "symbol grounding problem" (Harnad, 1990, 2002). This refers to the fact that a symbol is not just a free-floating token, ready to be manipulated according to abstract rules, but a token with a meaning, i.e. a token that stands for some phenomenon or situation *outside* the symbol system. An RSS or language cannot specify its own interpretation, i.e. the process through which its symbols are to be connected to the phenomena they stand for. The reason is that these phenomena are not symbols themselves, and therefore they forever remain outside of the language that intends to describe them. This fundamental limitation has been called the "linguistic complementarity principle" by Löfgren (1992). A dictionary may explain one word in terms of other words, e.g. "darkness is the absence of light". Yet, in the end there will always be words, such as "light", "bitter", or "dog", that can only be really understood by seeing light, experiencing a bitter taste, or playing with an actual dog (Massé *et al.*, 2008). The process that connects the symbol to its non-symbolic meaning is called “grounding the symbol”. However, it is everything but clear how this process precisely takes place.

Previously, logicians had tried to evade the problem by using formal semantics, in which a correspondence is established between a symbolic representation and some “model” representing a possible world in which the symbolic expressions are supposed to be true (Copeland, 2002). But the model is just another abstract system of symbols. Thus, semantics, understood as a mapping from a formal symbol to a symbolically specified “object” that the symbol is supposed to denote, fails to solve the symbol-grounding problem. The real world does not consist of clearly delineated, persistent objects that exist in an objective, one-to-one correspondence with clearly delineated, persistent symbols. As demonstrated both by our subjective experience and by the failed experiments with robots programmed according to purely symbolic schemes, the real world is intrinsically complex, changing, ambiguous and subjective. Therefore, RSS always fall short in fully capturing its behavior.

The solution proposed to the symbol grounding problem is that true cognition is always *situated* and *embodied* (Steels and Brooks, 1995; Clark, 1998). That means that we cannot understand the world just by reasoning about symbolic representations. We also need a material body equipped with *sensors*, such as eyes and ears, that can perceive the external situation, and *actuators* or *effectors*, such as muscles and vocal chords, that can act and thus affect that situation. Perceptions and actions are connected by a feedback loop that runs through the environment (see Fig. 2):

- a situation is perceived;
- the perception is processed so as to make sense of it with respect to the agent’s goals
- this interpretation triggers an action to bring the situation closer to the goals,
- this action, together with external events, changes the situation;
- the new situation is perceived again, thus initiating a new round of the feedback loop

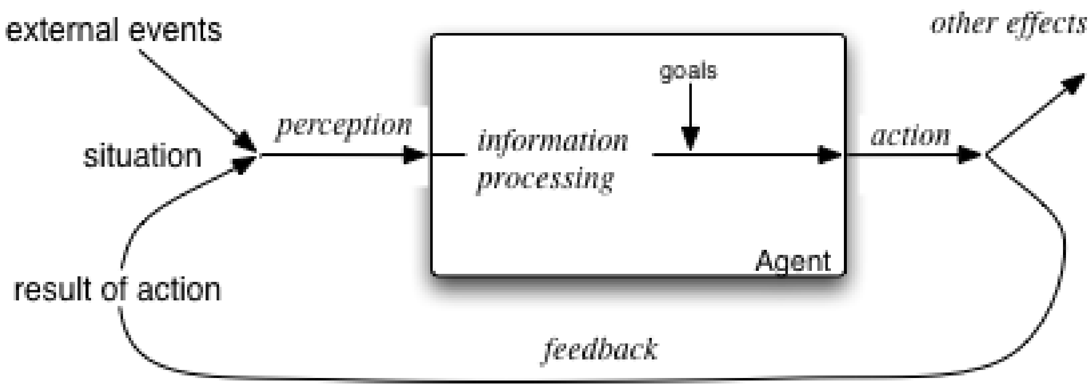

Fig. 2: a simple representation of the feedback loop in which an agent (typically a person) perceives the situation in its environment, processes that information while comparing it with its goals, and then acts to bring the situation closer to its goals. The result of that action is fed back to the perception, so that it can trigger subsequent actions to correct any new or remaining deviations from the goal.

Perceptions not only produce actions. They are compared with the expectations derived from previous perceptions and our knowledge of the world. If perception and expectation diverge, as typically happens in a non-trivial environment, a correction is to be made in the knowledge. Thus, the cognitive system constantly adapts its knowledge to *experience*, i.e. to its actual interactions with the world. These on-going interactions, which “enact” our cognitive processes (Stewart, Gapenne and Paolo, 2014), ground the cognitive system into the real world.

This system never settles into an objective representation of the world, in which each important object or aspect of the world would be accurately and permanently reflected by a corresponding symbol. This cognitive system therefore cannot be conceptualized as a collection of discrete, independent symbols. It is rather a fluid, evolving network of associations between different perceptions, interpretations and actions.

The most successful approach for modelling such associations can be found in so-called “connectionist” or “neural” networks (McLeod, Plunkett and Rolls, 1998). These are inspired by the activity of our brain’s neurons and their connecting synapses. Such networks learn and refine their knowledge by reinforcing connections that contributed to successful predictions or actions, while weakening the less useful connections. Therefore, their knowledge is not localized in any specific elements, concepts or symbols; it is *distributed* across the network of connections.

Such networks interpret perceptions, infer implications, and decide which actions to perform by means of a process of *spreading activation*. The neurons directly connected to sensors are “activated” by the incoming stimuli, such as light falling onto nerve cells in the retina at the back of the eye. Each neuron passes on its activation to the neurons it is connected with, with an amount proportional to the strength of the connection. The neurons in this next layer are activated proportionally to the sum of the activations they received from their incoming connections. They pass on this activation in the same way via their outgoing connections, to a subsequent layer of neurons. As this process is repeated, the pattern of activation is propagated step-by-step across the network, while changing with every step. The recently successful AI method of “deep learning” uses neural networks consisting of many layers in which the activation coming from perceptual input is processed into increasingly abstract patterns as it reaches subsequent layers (Schmidhuber, 2015). Eventually, the activation may reach effectors, which convert it into actions. This action affects the situation. Perception of the changed situation in turn activates the sensors, thus initiating a new round of spreading activation. Thus, activation continues circulating through sensors, brain, body and world in an endless feedback loop.

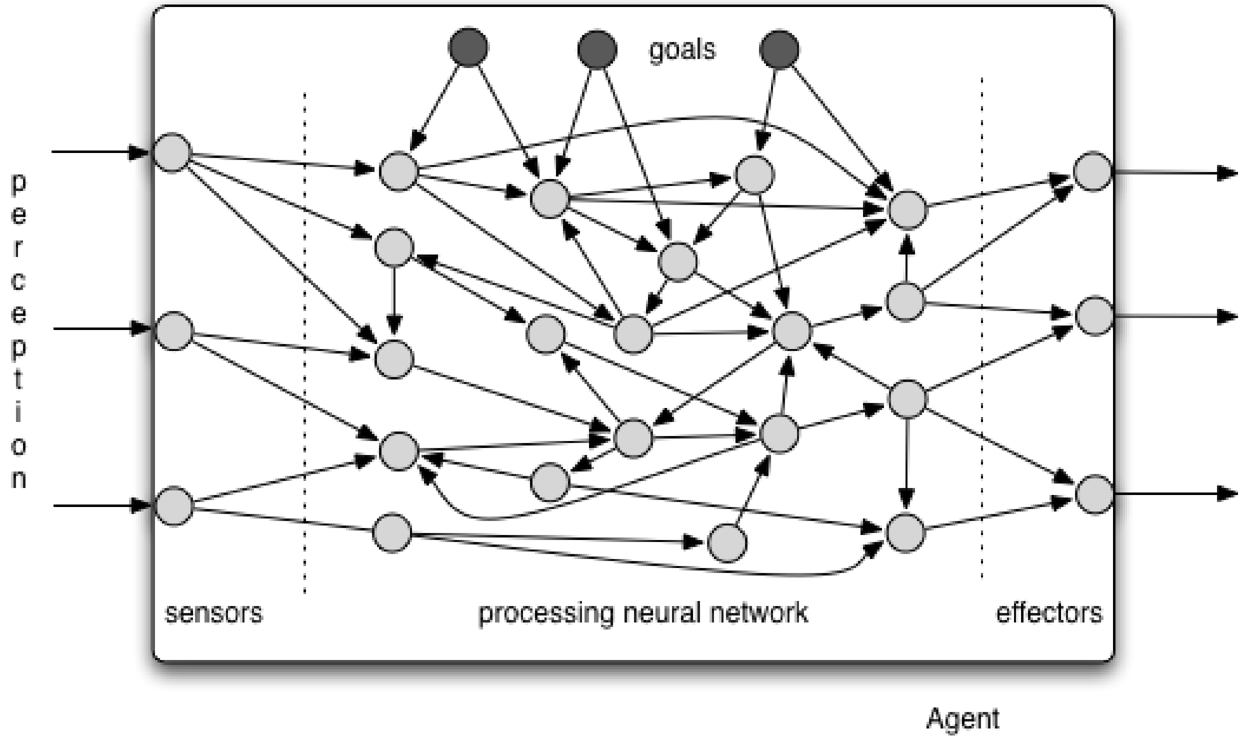

Fig. 3: a more detailed sketch of the neural processing happening within an agent. Sensors, neurons, effectors and goals are here represented by grey circles, playing the role of nodes in the network. Connections are represented by arrows linking one node to a next one. Perception of the outside world activates sensor nodes. These pass on their activation via their connections to subsequent layers of neuron nodes, until it finally reaches effector nodes that produce action.

This fluid, neural mechanism of cognitive processing, which appears to be the dominant one in the brain, has been called "subsymbolic" (Bechtel and Abrahamsen, 1991; Kelley, 2003). That means that it happens at a (mostly subconscious) level where no distinct symbols can be recognized as yet. The RSS-based cognition we know so well, with its grammatical and logical rules for manipulating symbols, turns out to be merely a superstructure grafted on top of this immensely complex and dynamic network of subliminal processes. From this perspective, rational thinking based on symbols is merely the tip of the iceberg, or perhaps more accurately, the visible shoot of a plant whose underground network of roots "grounds" these few leaves in a dark and rich soil without which they would not be able to survive and grow.

The Origin of Rational Symbol Systems

But where do those apparently simple and logical RSS come from? The emergence of symbol systems has been elucidated with the help of computer simulations of the origin of language (Steels, 1998, 2005). In these simulations, robotic or software agents collectively develop a shared vocabulary and system of rules to designate the phenomena they perceive through their sensors. The dynamics of these simulations has been confirmed by experiments with people, in which groups of individuals are stimulated to develop conventional labels and schemes in order to collaboratively tackle a problem for which they have no vocabulary available as yet (Garrod and Doherty, 1994; Fay *et al.*, 2010; Fusaroli and Tylén, 2012).

The general dynamics in these experiments and simulations is the following. Agents start out with some individual association between, on the one hand, a meaningful category of phenomena they can sense, on the other hand, a symbol or "name" to designate that category. When two agents meet, the one indicates a phenomenon that belongs to such a category and names it. When both agents agree about the designation, the association between name and category is reinforced in their mind. When they disagree, the association is weakened for the one that used it, and reinforced for the one that heard it, but did not recognize its designation. Thus, after such an encounter, their associations become a little more similar. Many such encounters take place between different agents using different names and pointing at different phenomena. Each time, initially different associations become more similar, while similar ones are strengthened. Eventually, all agents agree on all associations, thus becoming fully "aligned" in their symbol use: they have developed a consensual scheme for subdividing the reality they experience into a number of categories denoted by conventional symbols. Thus, an RSS has self-organized out of distributed communications between initially independent agents (Heylighen, 2013).

The underlying dynamics is characterized by positive feedback: because associations that are initially more common are more often reinforced, they become even more common. At the same time, they inhibit and eventually outcompete less common associations. This means that minimal differences in initial "popularity" are non-linearly amplified, leading to a "winner-takes-all" outcome in which a single symbol-category association erases all rival symbolizations. But in a different group of individuals, where by chance the initial distribution of associations was slightly different, other associations will win the competition, leading to a wholly different symbolic system, language, or culture (Axelrod, 1997; Henrich and Boyd, 1998).

The idiosyncrasy of such self-organized RSS is not limited to the "names" or symbolic tokens used to designate categories. The categories too undergo shifts and consolidations during the alignment process. Because there is no objective way to subdivide a continuous field of experience into distinct categories, different agents are likely to include different phenomena in a given category. For example, suppose that initially the agents have a somewhat different understanding of which category of colors the symbol "X" refers to. "X" may represent orangey red for the one, yellowish orange for another, and pale reddish brown for a third one. Whenever two agents encounter a color about which they agree on how to name it, that color becomes more firmly established in their mind as representative of the category denoted by the symbol. When one of them disagrees, the association between that color and the category is weakened for the other one. As illustrated by the study of Steels and Belpaeme (2005), this mutual alignment eventually settles on a category about which all agents now agree, e.g. "X" stands for orange.

Thus, the evolution of an RSS not only specifies the symbolic tokens, but also their standard meanings, i.e. the category of phenomena that a token stands for. This is a crucial step in the emergence of the typically human form of symbolic cognition.

Fig. 4: while the precise species of this insect-like creature (a type of praying mantis) can be unambiguously specified by its scientific name, its threatening yet elegant stance cannot be evoked merely by such RSS label (see text).

At the subsymbolic level, on the other hand, the meanings of perceived patterns are essentially subjective, context-dependent and transient. They only make sense in a particular situation for a particular individual, who uses particular associations, learned through a lifetime of particular experiences, to interpret that situation. The same phenomenon experienced by a different individual, at a different moment, or in a different context, is likely to be granted a different meaning (Heylighen, 1999). For example, the same insect-like creature (Fig. 4) may be perceived as an elegant appearance by one person, a disgusting bug by another one, or a rare specimen by a third one. An RSS largely eliminates these idiosyncratic associations and interpretations, and replaces them by an invariant, conventional scheme for categorizing and labeling phenomena. For example, the insect-like creature may be categorized as a praying mantis of the species *Archimantis latistyla.* This eliminates any ambiguity as to which type of creature was observed.

The emergence of a shared RSS simplifies and reduces the cognized phenomena to a relatively small set of invariant categories. This is because the process of communication that created the symbolic system needs to achieve a consensus among a large and disparate group of individuals. It can only do so by restricting itself to the relatively coarse and salient distinctions that everyone in the group is capable of making, while disregarding more subtle differences that certain more sensitive individuals may be able to perceive. The collection of symbols also needs to be small and simple enough so that most people in the group can understand and memorize them, and so that there are sufficient on-going communications that use these symbols to regularly reinforce the association between symbol and meaning. Symbols that do not fulfill these conditions will sooner or later be forgotten or lose their original meaning—a common event in the evolution of language. The advantage of such selective retention is that the symbols most likely to survive the process are those that represent the most stable, useful and clearly delineated categories—such as body parts, biological species or common types of household objects. Learning those through education into the RSS provides the mind with ready-made conceptual tools that can be used in a wide variety of circumstances to describe and reason about common situations.

This creation through communication of a system of shared categories has been called the *social construction of reality* (Berger & Luckmann 1966): each social system will develop its own symbols, categories and rules to represent reality. In the terminology of AI and ICT, such a consensual system of symbolic distinctions is called an *ontology* (Livet, 2012; Steels, 1998)—originally a philosophical term for the study of the basic elements of reality. Thus, an RSS, by reducing the continuous and ever-changing field of experience to a finite set of discrete and persistent categories, imposes a particular view of what reality consists of. Because there is no objective way to capture the whole of reality, this restricted view is always to an important degree (inter)subjective, i.e. dependent on the local culture. Therefore, different communities that have settled on different categorizations will not really be able to communicate about their differences.

Translation is only possible if the symbols of the first system can be mapped precisely onto those of the other system while keeping the same interpretation. However, that assumes that the two RSS would make exactly the same distinctions. This is in general not the case. For example, the meaning of the French verb "aimer" roughly covers the same meaning as the two distinct English verbs "to love" and "to like". Therefore, the expression "j'aime" can be translated as either "I love" or "I like", and only a deeper understanding of the context may indicate which translation is the right one. Such intrinsic ambiguities in part explain why symbolic AI failed to achieve automatic translation. AI researchers initially assumed that the words of all languages could be mapped onto universal semantic categories—an underlying "language of thought". However, they failed to uncover such a fundamental, culture-independent ontology. Therefore their attempts at translation eventually had to switch to more pragmatic, "subsymbolic" methods. Here, the most likely meaning of a word or phrase is inferred indirectly from the words in the context and the statistically derived associations between these words and phrases rather than from a dictionary definition (Heylighen, 2001).

More generally, aspects of reality for which there are no corresponding symbols or categories in the given RSS simply cannot be described, communicated, stored or even rationally reflected about. Thus, while the use of rational symbol systems has given humanity the power to represent, register and reason, it has done so only in a highly limited, idiosyncratic manner.

Transcending RSS limitations

Metasystem transitions

Cybernetics describes living organisms as goal-directed systems (Heylighen and Joslyn, 2003). Such systems achieve their goals (desired situations) by counteracting any deviation from these goals as perceived through their sensors (Figs. 2 and 3). In that way, they control their situation so as to come as near as possible to the situations they desire. To succeed in that, they need knowledge in the form of rules that tell them which action is appropriate for which situation. Given adequate sensors, actions and rules, they can adapt to a wide variety of situations, and deal with the most common problems that are likely to occur. However, this adaptivity in behavior is not reflected in the rules that tell the system how to adapt. In principle, the agent could become even more adaptive, and thus intelligent, if its rules too could adapt to the circumstances. But that requires a higher-order system of "meta-rules" that control the underlying rules, so that they can change these rules to better fit the perceived situation. Such a higher-order system is called a *metasystem*. Its emergence over the course of evolution is called a *metasystem transition* (Turchin, 1977, 1995; Heylighen, 1995).

The neural or subsymbolic level of cognition is already a metasystem with respect to the more primitive level of *reflex* actions. Reflexes, like pulling your hand away from a hot surface or breathing in when your body needs oxygen, are innate and do not need to be learned. They use pre-existing, hard-wired rules that are useful in a wide range of circumstances. But, like for all rules, there are cases where reflex rules are inappropriate—e.g. when trying to breathe in the presence of poisonous gas. In such circumstances, it would be good if they could be overruled and replaced by more appropriate rules. Moreover, most situations are too complex to be covered by the relatively small number of such innate reflexes. That is why evolution gave organisms the ability to adapt existing rules and to develop new ones through the process of *learning*. This happens through the mechanism of neural reinforcement that we sketched earlier, in which connections that worked well are strengthened, while those that did not work are weakened. Thus, the emergence of reinforcement learning constitutes a metasystem transition (Turchin, 1977; Heylighen, 1991).

The emergence of RSS constitutes a subsequent metasystem transition. The problem it solves is that the reinforcement rules that govern learning only work on the basis of experience. The organism needs to actually experience a situation in order to learn what works and what does not work in that situation. As we saw, symbolic cognition overcomes that limitation by using combinations of symbols to conceive of situations that have never been experienced, while using logic to infer what would or should happen in such a situation. In this way, symbolic cognition allows us to formulate rules applicable to situations that have not been encountered yet, thus widely extending the range of situations to which we can adapt. But RSS too are limited, because their symbols and rules are still restricted, rigid, and dependent on idiosyncratic processes of social construction that we do not control.

The question then is whether we can overcome these limitations as well. This would initiate a next metasystem transition that would allow concepts and symbolic rules to be freely extended and adapted to the circumstances. Science, art, philosophy and spirituality all have approached that problem, proposing different methods to go beyond these limits. In the following sections, we will briefly review the solutions that these different cultural movements have proposed. These solutions developed roughly in two stages (Heylighen, 1984): a first, "classical" one in which a shortcoming was noted and one or more apparent solutions were proposed, and a second, "revolutionary" one, typically starting in the first half of the 20th century, in which the remaining limitations of any such solutions became clear and a more radical exploration and experimentation was initiated that questioned the RSS to its core. In the last section, we will explore in how far 21st century ICT may provide the medium that allows these "post-RSS" developments to become coordinated and integrated into an emerging "Global Brain" that would provide a true metasystem with respect to the RSS.

Science: formalizing and operationalizing concepts and models

Science can be characterized as an attempt to develop knowledge without the simplifications, subjectivity and ambiguity that characterizes everyday knowledge as expressed in natural language. Science in principle strives towards an as accurate, objective and complete as possible representation of reality. To achieve this, science uses the methods of formalization and operationalization (Heylighen, 1999).

Formalization means that symbols, rules and expressions are defined explicitly, like in mathematics or programming languages, so that their interpretation does not depend on the person, context or situation. For example, the symbols "0" (zero) and "+" (addition) are defined by the axioms characterizing the system of numbers. This includes the axiom stating that: $n + 0 = n$, for any number n. Therefore, such formal symbols are understood by everyone in exactly the same way, without ambiguity (Fig. 5). Another advantage of formalization is that in principle you can make symbols have any meaning you want, while maintaining a logically coherent system, simply by characterizing them by the right axioms and definitions. This allows mathematicians to conceive and investigate abstract algebras and hyperdimensional geometries that do not seem to have any counterpart in physical reality (although one may be found eventually).

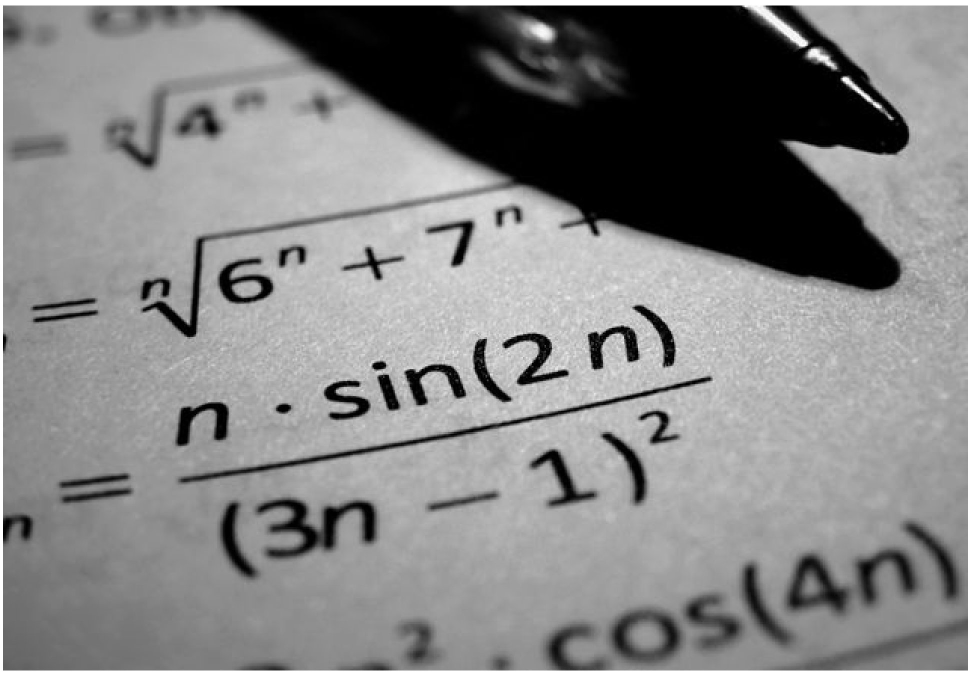

Fig.5: science makes knowledge more reliable and objective by formalization, i.e. by representing it with symbols that have a precise, unambiguous definition in terms of other symbols.

Still, formalization can only define symbols in terms of other symbols. It thus lacks a way to ground symbols in external reality. The latter method is provided by *operationalization*. This is the definition of a scientific concept or category by means of the concrete physical operation needed to establish whether some phenomenon belongs to that category. For example, to determine what an object's weight is you can put it on a scale and read off the value. Similarly, you can operationalize the concept of "introversion" by developing a psychological test that dependably discriminates the more introverted people from the more extroverted ones.

Formalization and operationalization make it possible to reliably establish the correctness of hypotheses. In a completely formal theory, the truth or falsity of a proposition can in principle be established by means of a proof, which derives the proposition (or its negation) from axioms by applying logical rules of inference. In a theory with operationally defined properties, the truth of a proposition involving such properties can be established by performing the operations and checking whether they produce the predicted results. For example, you can test the theory that says that objects orbiting in space are weightless by putting the same object on a scale on Earth and in a space station, and comparing the results. Thus, the great power of science is that it provides a method to reliably test hypotheses, so that it can eliminate the bad ones and keep the better ones. This provides it with a remarkably effective engine of progress, facilitating the discovery of dependable categories and rules.

Up until the beginning of the 20th century, it was assumed that, using these methods, science would eventually be able to develop a complete, perfectly accurate, and objective representation of reality, in which everything would be determined by formally defined and operationally specified symbols and rules. However, in the first half of the 20th century an array of "limitation principles" was discovered (Barrow, 1998; Yanofsky, 2013) that shattered this dream. The limits of formalization were established by the *theorem of Gödel* (Gödel, 1992; Franzén, 2005), which demonstrates that even in a completely formal RSS we will never be able to derive the truth or falsity of certain propositions from the axioms of the system. A similar limitation was formulated as the *halting problem*: in general we cannot determine whether a formally specified computer program will come to some conclusion or continue to run without end (Burkholder, 1987). The limits of operationalization were made clear by the *Heisenberg uncertainty principle* in quantum mechanics (Hilgevoord and Uffink, 2001), and the *butterfly effect* in complex systems (Hilborn, 2004). These principles show that we will never be able to accurately establish the state of some external phenomenon because the operation of observation necessarily perturbs the phenomenon in some unpredictable manner, while even the tiniest perturbations can have huge effects that invalidate any prediction derived from the initial observation.

The conclusion is that science cannot hope for any complete, observer-independent description of reality. Formal and operational representations are intrinsically limited (Heylighen, 1999). The linguistic complementarity principle, which can be seen as a generalization of these limitations (Löfgren, 1992), cannot be evaded. While the scientific method allow scientists to overcome many sources of subjectivity, ambiguity and inaccuracy, they must accept that in the end their theories are merely *models*, i.e. simplified representations that reflect only limited aspects of an infinitely complex reality (Bailer-Jones, 2009). Including too many aspects makes the model too complex to process. Therefore, different aspects or problems concerning the same situation require different models—without any one of them being the "true" model of the situation. In practice, then, models are chosen by convention, because working scientists agree that they happen to work well for a given type of problem.

This more pragmatic attitude, which was called "conventionalism" by the mathematician Poincaré (Folina, 2010), has slowly diffused throughout most scientific disciplines. It is perhaps developed most explicitly in the sciences studying complex systems—such as societies, organisms, or the brain—whose subject is so complex and dynamic that any attempt at building a complete representation is anyway doomed to failure. These sciences have developed a variety of broad-purpose modeling techniques, including systems analysis, systems dynamics, network analysis, dynamic systems and multi-agent simulations. Depending on the needs of the situation and the available computing capacities, such models are easily adjusted, incorporating more or fewer features.

For practical purposes of prediction, these methods are often combined with methods developed in AI and ICT for machine learning, knowledge discovery or data mining, i.e. for the extraction of recurrent patterns from huge amounts of data. The difference is that the models are proposed by the scientists and then adjusted to the empirical data, while the machine learning methods start from the data, while trying to fit these into common models. Machine learning methods include neural networks and deep learning, genetic algorithms, statistical analysis, and probabilistic induction and abduction. Rather than using the categories and rules of a universal RSS to build the supposedly true representation of a phenomenon, such scientific modeling and data mining techniques help us to find new (fuzzy) categories and rules tailored to the specific problem domain.

In that way, they externalize our ability of *inducing* concepts and rules from experience. The human brain has an in-built capability to recognize regularities in myriads of scattered observations, and thus infer higher-level, abstract categories and rules. However, this capability functions almost exclusively at the subsymbolic level, where it was hitherto inaccessible for communication, registration, reflection or manipulation. Contemporary science and ICT are learning to externalize this capability in mathematical models and computer programs, thus making it accessible for practical problem solving and continuing improvement. But to achieve this, they had to give up the classic philosophical assumption that there is a permanent and objective one-to-one correspondence between the symbols of a representation and the objects in the real world. Instead, the (mostly mathematical) symbolizations they use represent the ever-changing connections, processes and interactions between nodes of activity. In that way, their dynamics and organization increasingly resemble the one of the neural networks in our brain.

Philosophy: from answering questions to questioning answers

Philosophy, like science, can be seen as an attempt to go beyond the superficiality, ambiguity and subjectivity of RSS in trying to understand reality. Rather than formalization and operationalization, its main tool is critical reflection. Initially, this reflection was directed at filling the gaps in the RSS ontology: answering fundamental questions suggested, but left unanswered, by standard categories and rules. For example, according to the rules of causality, every event must be produced by a prior cause. But that leaves open the question of what caused the universe to come into being, given that this initial cause by definition must have been outside the universe. This led Aristotle (Fig. 5) to postulate a "Prime Mover" or "Uncaused Cause" that created the universe (Bodnar, 2018).

As philosophy became separate from science, the questions it focused on were those that did not allow any formal or operational approach, i.e. that could not be resolved through observation, computation or demonstration. These were initially mostly problems of metaphysics or ontology, focusing on the ultimate categories and causes behind the phenomena we observe. For example, a long-standing division in philosophy is the one between materialists, for whom all phenomena are ultimately reducible to material particles, and idealists, for whom all phenomena are constituted by abstract ideas or forms.

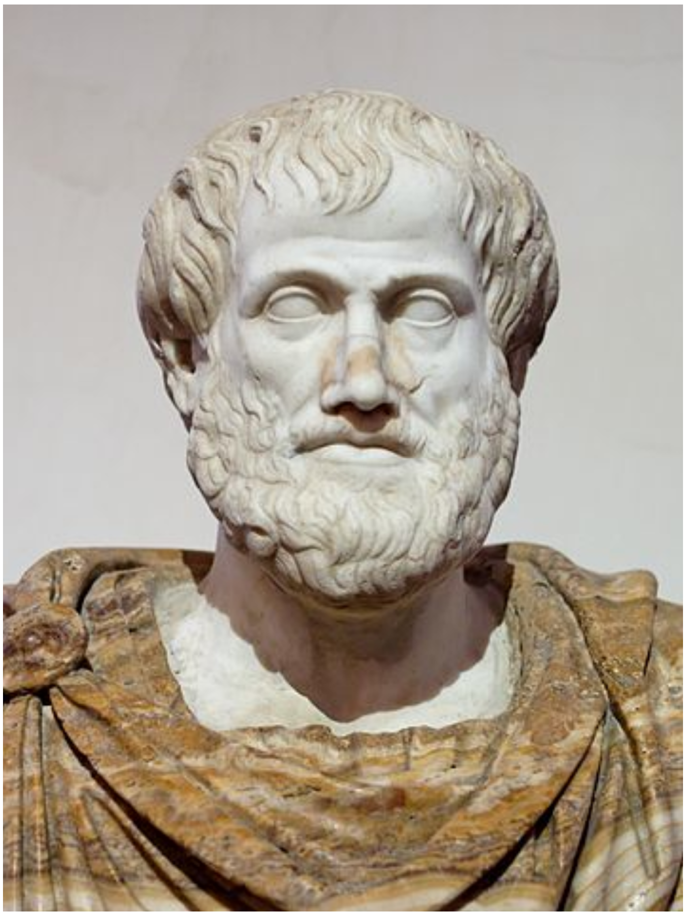

Fig. 6: Bust of Aristotle, the Greek philosopher who postulated a "Prime Mover" to explain the origin of the universe

However, as philosophy evolved, the number of philosophical concepts and systems proposed to answer such questions merely seemed to multiply. There was not any clear sign of progress, in the sense that newer answers would in some objective way be "better" than older ones, or that philosophers would reach some degree of consensus about the ultimate constituents of reality. Thus, it transpired that philosophical systems, in spite of their ambitions and the deep reflection on which they are based, are just as incomplete, subjective, and ambiguous as any other symbolic systems. This observation was perhaps formulated most forcefully in the early 20th century by Wittgenstein (2010), who noted that all such systems (and RSS in general) are merely “language games” that are intrinsically limited in what they can express (Black, 1979). This made Wittgenstein conclude that metaphysical discussion is merely a waste of time, as expressed in his famous quote: “Whereof one cannot speak, thereof one must be silent”.

This critical stance was highly influential. As a result, philosophy started to focus almost exclusively on a critical analysis of existing RSS, emphasizing limitations and raising further questions rather than proposing solutions. Thus, the focus moved from ontology to epistemology, i.e. the study of how we can gather knowledge about reality rather than the study of reality itself. This critical examination spread to encompass language, morality, religion, society, science, culture, media, technology, art… i.e. the symbolic systems that cover all domains of human life. Two major traditions are conventionally distinguished in such philosophical examination: the Anglo-Saxon, “analytic” approach (Glock, 2008), and the Continental, “hermeneutic” or “deconstructive” approach (West, 2000).

The analytic tradition tries to remain close to the formal and operational methods of science by introducing explicit categories and distinctions, and using logical reasoning to draw conclusions from “thought experiments” based on such proposed distinctions. However, since the categories these philosophers use by definition fall outside the categories that have been successfully formalized and operationalized in science, the conclusions drawn from such “counter-factual” reasoning rarely reach results that are concrete and consensual enough to make an impact outside of philosophy. At best, they help scientists to more precisely formulate the questions they are addressing, e.g. in modeling, cognitive science and AI.

The Continental tradition has moved farther away from science and logic and closer to literature and art, formulating its insights in often poetic but typically vague and ambiguous narratives. Its emphasis is on “deconstructing” the hidden assumptions and biases behind common ways of thinking and acting. Thus, the recent Continental tradition, especially in its “post-modern” guise, has been reminding us of the subjectivity, culture-dependence and implicit biases of common RSS, while inspiring us to go beyond such reductionist approaches by conceiving the world as a dynamic and holistic network of interactions (Heylighen, Cilliers and Gershenson, 2007).

Art: evoking experiences and inspiring insights

I will here consider art broadly to include not just the traditional domains of painting and sculpture, but poetry, literature, music, cinema, performance, theatre, dance and any other “creative” forms of personal expression. Art, like science, can be seen as an attempt to transcend the limitations of symbolic communication and cognition as rooted in verbal language. However, while science strives to eliminate the subjective aspects of symbolic representation, art in a way does the opposite: providing access to people’s highly subjective, personal meanings and experiences, which are too subtle and idiosyncratic to express in a system of discrete symbols and rules. The subtle combination of curiosity, beauty, fear and disgust elicited by the sight of a praying mantis (Fig. 4) cannot be expressed by its scientific name *Archimantis latistyla*. However, a drawing, animation or poetic description of that creature's threatening but elegant stance might. Art in a sense bypasses the rational, symbolic level of cognition and directly addresses the subsymbolic level of associations, feelings, and experiences, where it elicits the kind of intuitive meanings that cannot really be expressed in words.

To achieve that, art makes use of two methods of signification: icons and metaphors. In semiotics, the science of signs (Merrell, 2001; Hawkes, 2003), an *icon* is defined as a sign that conveys a particular meaning by its resemblance to the phenomenon it refers to. For example, a portrait refers to the person being depicted by its similarity in shape to the person's facial features. Someone looking at the portrait will be immediately reminded of the person it represents, and therefore of that person's character, typical expressions, and walk of life—all features that are impossible to fully capture in words. Yet, a portrait is more than a photograph—even though the latter arguably captures the person's facial features more accurately and in more detail. Just like in the case of scientific models, the strength of an artistic representation lies in simplification and selectivity, in expressing the features that are most meaningful or important to convey, while not clouding the picture with unnecessary details. What is important in this case is determined by the artist and the subject, without any pretense at objectivity, but rather at expressing the, for the artist, most intuitively or emotionally significant aspects of the subject. These are also the ones that are most likely to move or inspire the viewer of the artwork. Thus, art ideally manages to register and communicate the subjective, emotional, intuitive aspects of some idea or phenomenon—which are ignored by RSS—by evoking the corresponding feelings in the audience.

Such subjective selectivity of representation is even more outspoken when using metaphors. Like an icon, a *metaphor* uses the similarity between a symbol and what it stands for, but this time the similarity is to be found not with the shape of the symbolic token (the signifier), but with the phenomenon it refers to (the signified). For example, if you describe a certain person as a "wolf" you refer to a certain similarity between the characteristics of that person and the characteristics of the animal that is conventionally called "wolf". Like in the case of an icon, this similarity evokes subjective feelings and associations (such as hunger, greed or brutality) that are not evoked by the RSS labeling of that person. These feelings are intrinsically ambiguous and context-dependent, because the metaphor does not specify which characteristics are shared between the thing being literally referred to (the wolf) and the thing metaphorically referred to (the person). Obviously, that person does not have a grey fur and a long snout with sharp canines, but he might have a large appetite or an aggressive disposition. When used in the context of eating, the metaphor may refer to the appetite; when used in the context of interpersonal relationships, it may refer to being aggressive. Like in the case of iconic art, the skill of the writer or poet describing a phenomenon lies in selecting the combination of metaphorical elements and context-elements that would evoke the most pertinent feelings in the reader.

The use of icons and metaphors allows art to create symbolic representations that reflect our subsymbolic experiences and intuitions better than RSS do. Artistic representations convey not just the category to which a phenomenon belongs, but the connotations and affective meanings associated with that phenomenon. This incites the mind to make further associations, by letting the neural activation spread from the evoked meanings to more indirectly related meanings that put the phenomenon in a broader perspective or that suggest novel aspects. For example, a person who has read a poem, seen a painting or watched a movie may be inspired by this work of art to develop a broader, intuitive understanding of the situation depicted and to see it in a new light. This cognitive processing plays a role similar to the rational, logical inference supported by an RSS representation in extending understanding, but now at the more fluid, felt, "subsymbolic" level. Thus, artistic representation not only communicates felt meanings, it also helps us to process that meaning into novel insights.

However, this latter function is as yet much less developed than the logical processing enabled by rational symbol systems. Initially, art seems to have been content to represent easily recognizable meanings by using traditional formats of expression, such as stories, figurative paintings, or songs. But the beginning of the 20th century witnessed a revolution in which art began to radically experiment with new representations, new meanings and new ways of interaction with the public. This put into question everything that had come before. This revolution was exemplified by a wave of "avant-garde" or "experimental" movements and approaches (Poggioli, 1981; Wood, 2004), including cubism, Dadaism, surrealism, abstract expressionism, avant-garde theater, atonal music, electro-acoustic music, video art, performance art, installations, mixed-media and multi-media, computer-generated art, and many more.

What these movements have in common is a desire to experiment with novel forms of expressions and novel combinations of existing forms, thus transcending the boundaries between different disciplines, genres and media. For example, a contemporary theater piece or performance may exhibit actors that execute movements and dialogues in a partially scripted, partially improvised manner, a background screen on which video-taped or computer-generated movies or pictures are shown, an audio background of music and electronically generated sounds that accentuate or suppress the actors' speech, physical objects that are assembled by the actors in different configurations or that are programmed to react in certain ways, and a public that interacts with the events on stage by proposing additional activities via their smartphone.

Such experimental pieces and performances are intended to question the assumptions and expectations of the conventional RSS, just like critical philosophy does. Their bringing together of disparate but evocative representations ideally elicits novel insights—in the artist, the performer, or the audience. By putting together meaningful symbols that are normally never put together, such art incites the creation of new associations and combinations in the brain, thus stimulating imagination, creativity and discovery.

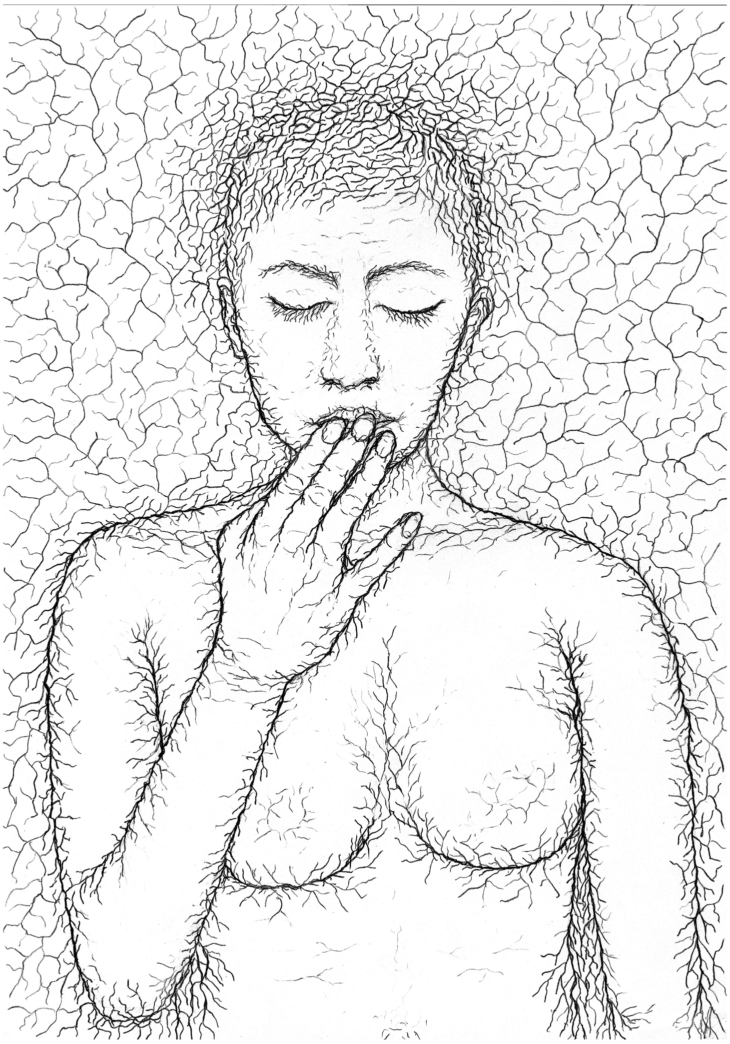

Fig. 7: an experimental work of art (here a drawing by the author) tries to elicit new meanings and associations by combining evocative representations (here depicting a torso of a woman together with a pattern of cracks reminiscent of drying clay) that are normally never seen together

Such unconstrained combination was perhaps explored most systematically by the surrealists (Matthews, 1986), whose painters produced unforgettable images such as melting clocks, bowler-hatted men raining from the sky, or an apple that fills a room. A favorite surrealist game illustrates the method at its simplest. Different people each in turn add some element to an emerging sentence or drawing, albeit without seeing what the others have previously contributed. This leads to apparently absurd results, such as “Le cadavre exquis boira le vin nouveau” (“The exquisite corpse shall drink the new wine”), the sentence that gave the game its name of “Cadavre Exquis” (Adamowicz, 1998). From the surrealist perspective, such juxtaposition of logically unrelated elements provides a window into the subconscious—which is also the realm of dreams—by inciting the mind to find intuitive, non-rational connections between these elements. Thus, surrealism and related art movements that use “impossible” combinations of symbols (Fig. 7) can be seen as searching for “surrealities” that go beyond the socially constructed reality of the conventional RSS.

More mainstream art—as exemplified perhaps by Hollywood movies, horror stories, and pictures of maidens and unicorns—is not as radical in its exploration of alternate realities and experimentation with formats and symbols. Still, it has assimilated the most successful results of previous experiments, incorporating plenty of surreal fantasies, absurd humor, and computer-generated imagery. The popularity of genres such as thrillers, science fiction, and fantasy testifies to the fact that an underlying message of even commercial art is that reality is not what it seems, and that the situation may be very different from what the reigning RSS has taught us. Concerts, performances and movies integrate a variety of sophisticated media, technologies and methods that together create an all-encompassing spectacle that speaks to the major senses and emotions. Thus, art has fully embraced ICT and other technologies to more powerfully convey experiences and explore creative processes. The recent approach of ArtScience (Edwards, 2008; Root-Bernstein *et al.*, 2011; Siler, 2011) even attempts to synthesize the more playful, serendipitous artistic methods of experimentation and exploration with the more controlled scientific ones—e.g. by letting computer programs process astronomical data into poetic text (Petrovic, 2018), or trying to recreate the origin of life by subjecting mixtures of molecules in transparent vessels to electrical discharges.

Spirituality: controlling attention and expanding consciousness

Next to science, philosophy, and art, there is a final domain of human culture that has systematically attempted to overcome the limitations of rational, symbolic cognition. This not very clearly defined domain has been called *mysticism*, *spirituality*, or *consciousness expansion* (Glenn, 1989; Deikman, 2000; Harris, 2014). It is exemplified by practices of meditation, mindfulness and the use of psychedelic drugs. Such techniques help people to bypass the rigid distinctions imposed by the RSS, and thus achieve a more direct access to experience and the reality behind it. While such spiritual practices are often associated with religion, as a supposed means to attain communion with deities, they are in a sense antithetical to the more common understanding of religion as a collection of articles of faith that are to be accepted without questioning (Harris, 2014). Whereas religious faith is supposed to be "blind", spiritual practice is supposed to make its practitioners see more clearly, by opening their "doors of perception" (Huxley, 2009) and thus helping them to achieve "enlightenment".

The tradition is ancient, going back to shamans that used hallucinogenic mushrooms, fasting, sweat lodges, and other intense experiences to induce a trance-like state that would bring them into contact with the "spirit world" (Stutley, 2003). Some of these practices and the accompanying philosophies were systematically developed in Eastern spiritual traditions—in particular Buddhism, Hinduism, Taoism and Sufism (Kohn, 2008; Smith and Marranca, 2009). Here, the focus is on learning to control one's desires, thoughts, feelings, and most generally consciousness. This allows practitioners to detach themselves from the conventional "reality" of the RSS, while getting in touch with a wider or deeper reality behind it. The most common method is *meditation*: unwavering concentration on a single idea or sensation (Fig. 8). This is intrinsically very difficult because attention naturally wanders towards ever-new feelings and thoughts. When concentration is maintained sufficiently long, this may result in a "mystical experience", in which the subject loses the sense of being a separate individual here and now, and instead feels like becoming one with the universe—the so-called "oceanic feeling" (Parsons, 1999).

Fig. 8: a common spiritual method to detach oneself from the conventional, RSS-mediated way of experiencing reality is *meditation*: unwavering concentration on a single idea or sensation

Deikman (2000; 1966) has interpreted such an experience as a “deautomatization” of the neural processes in our brain. Normally, activation is propagating automatically along the strongest, most frequently used connections between neurons, thus helping us to make sense of the present situation within our standard picture of reality. The continuous refocusing of attention on the same neural region, however, creates an “overload” of these strong connections, so that they stop performing their habitual function. You can demonstrate this effect to yourself by continually repeating the same word—say “table”. After a while, the word will lose all of its meaning, turning into a mere sound rather than a symbol with a signification. If the brain region that is overloaded includes the neural systems that monitor who, when or where you are, you may even lose your sense of self and instead feel like a timeless emptiness.

The overload of strong, automated connections allows activation to spread along weaker, less common connections that reach more indirectly associated meanings. The result is a “fresh” experience, where the same phenomenon is perceived in a different light—typically more vividly and in more detail, while triggering novel sensations and creative insights. Related mystical experiences can be produced by psychedelic drugs (Sessa, 2013), such as LSD or psylocybin. These appear to facilitate the flow of activation along more diverse and unusual connections (Schartner *et al.*, 2017), thus alerting people to different aspects of reality. Unlike meditation, however, the effect of such hallucinogenic drugs is difficult to predict and control, and can have dangerous side-effects, such as triggering "bad trips", psychotic episodes or losing touch with reality.

Korzybski's (1958) theory of General Semantics explains why consciousness expanding techniques are useful to help us overcome the limitations of symbolic cognition. General Semantics notes that our understanding of the phenomena we experience is built up in subsequent stages of neurophysiological processing: from the outside phenomenon, via our sensory experience, to an increasingly abstract representation of that experience, concluding with a particular word or symbol that supposedly captures the essence or meaning of that phenomenon. The problem is that we tend to identify that highly simplified and rigidified symbolic representation with the phenomenon itself—a fundamental mistake summarized by Korzybski's famous aphorism: "the map is not the territory". To counter the "insanity" that arises from this systematically repeated error, Korzybski (1958) proposes a number of psychological techniques that should make us again more conscious of the previous stages in our cognitive processing, when the phenomenon was not yet reduced to its symbolic caricature.

Numerous other psychological, therapeutic, meditative, physical (such as yoga or tai-chi), psychedelic and other methods have been proposed to expand consciousness and liberate people from the too rigid thoughts, categories and perceptions induced by their RSS (Harris, 2014). Perhaps the simplest and presently most popular approach is *mindfulness* (Bishop *et al.*, 2004). Here people are taught to become aware of all the fleeting sensations and feelings that pass through their body and mind at any moment, focusing on these subsymbolic experiences rather than on the verbal, symbolic reflections and ruminations that most often dominate their consciousness. This practice has been shown to increase mental health and well-being in a wide variety of situations (Ludwig and Kabat-Zinn, 2008).

Like science and art, consciousness-enhancing techniques can be supported by technology. Researchers are increasingly concluding that consciousness is not just determined by the brain but by the whole of our interactions with the outside world (Noë, 2010). This observation parallels the "situated and embodied" and "enactive" views of cognition (Clark, 2008; Stewart, Gapenne and Paolo, 2014), which note that mental content must be grounded in such interactions to be meaningful. That implies that external tools that support such interaction also enhance cognition and consciousness. The simplest such tools, such as microphones, infrared goggles, or telescopes, amplify our powers of perceptions. Others, such as walking sticks, hammers or cars, amplify our powers of action. Research has shown that if such tools immediately and dependably extend our powers of interaction we will start to experience them as part of our own body, and eventually consciousness (Clark, 2008). For example, blind people who have light sensors implanted that stimulate a fine-grained array of nerves in their tongue, eventually learn to associate the stimuli experienced on their tongue with outside patterns in such a direct way that it feels to them as if they are "seeing" these patterns inside their mind—while ignoring the implanted hardware (Kupers and Ptito, 2004).

Virtual reality (VR) is a technique in which a computer-generated three-dimensional environment is presented directly before the eyes via displays in a special headset. When the person wearing the VR-headset moves, the image in the display moves accordingly, so that it feels as if the person is actually moving inside the virtual world. As a result, the consciousness of the real world vanishes, to be replaced by a felt presence inside the virtual reality. This technique can obviously be used to create “psychedelic”, “out-of-body” experiences, in which the person e.g. feels as if flying in space and observing the Earth and the Moon. But it can also be used to tackle particular problems of “insanity”, in which inappropriate interpretations have become rigidified. For example, psychotherapists have started to use such virtual experiences to teach traumatized patients to better deal with the objects of their fears (Sanchez-Vives and Slater, 2005; Bohil, Alicea and Biocca, 2011).

Augmented reality is a technology that uses special glasses to superpose a computer-generated representation on top of an image of the physical reality, thus enhancing our unaided perception (Van Krevelen and Poelman, 2010). For example, while watching an ocean-liner from the outside, the superposed image might show you the distribution of the rooms and decks inside the ship, while informing you about its destination, and the year it was built.

The next technological breakthrough may well be a direct brain-computer interface, in which our brain waves or other neural signals are registered by sensors and interpreted by a computer program (He *et al.*, 2013). This would allow us to manipulate virtual representations or steer physical processes simply by thinking about them. The futurist Glenn (1989) has extrapolated such ever more intimate connections between brain and world to what he calls “Conscious-Technology” or a merger between technology and spirituality. But this merger is likely to also encompass science, art and philosophy. Let us then examine how ICT may be able to integrate all the different tools and techniques we have surveyed.

Towards a Global Brain

We have reviewed how symbols form the foundation for the typically human level of intelligence, which allows us to communicate, register and accumulate knowledge; to reason about that knowledge; and to conceive of situations we have never experienced before. This form of cognition is based on *rational symbol systems*, as exemplified by verbal language. These RSS allow us to combine discrete symbols into meaningful expressions that can be processed into further expressions. Such symbols behave like standard semantic components or building blocks, which can be stored, transmitted, combined and manipulated according to logical rules without thereby losing their meaning. That is because they provide stable externalizations of otherwise fleeting mental content.

However, this strength of RSS is also their weakness: meanings that are fluid, subjective, interdependent, dynamic or contextual cannot be expressed by such independent, static symbols. The failed attempts to build a human-level artificial intelligence based on RSS have reminded us that the bulk of human cognition still resides at the subsymbolic level of intuition and subjective experience. Symbols and rules on their own are not sufficient: they must be grounded in an immensely complex and dynamic network of interactions. These are carried by our body and brain, using billions of neurons, synapses, sensors, effectors, and their extensions into the outside world. Recent advances in science, technology, philosophy, art, and spirituality are giving us an increasingly broad and deep grasp of these subsymbolic networks and processes.

Yet, these approaches are themselves symbolic enterprises: they propose external representations that can be communicated, stored, and manipulated. The apparent paradox can be resolved as follows: these representations are actually *metarepresentations* (Heylighen, 1988, 1990). They represent not given categories or distinctions, but the processes that create such symbolic distinctions. For example, the "deep learning" neural networks of contemporary AI (Schmidhuber, 2015) are not reasoning about specified symbols or categories: they are inducing new categories by searching for recurrent patterns in immense arrays of low-level data collected by various sensors. For example, when such a program is trained with thousands of photos of pets harvested from the web, it will eventually learn to classify them in high-level categories, which we may recognize as cats, dogs, mice, parrots, goldfish, etc., but which were initially unknown to the program. This means that we have succeeded in externalizing some of the subsymbolic capabilities of our brain for discovering new patterns.

A deep-learning program is still a symbolic representation of some of our neural mechanisms of pattern recognition and sense-making. Thus, followers of Newell and Simon (1976) could argue that this new form of AI has vindicated their "physical symbol system hypothesis" that sees symbols as the necessary and sufficient condition for intelligence (Nilsson, 2007). Nevertheless, the organization of such a connectionist program is completely different from the one of a RSS: its "symbols" are no longer static, independent and governed by formal rules. Instead, they are constantly adapting, interconnected, and grounded in an ongoing interaction with the real world. Thus, the program is really a metarepresentation: a system that represents the processes that generate specific representations, such as the symbolic category of "cats", but that could equally learn to generate an infinite variety of other representations.

The shift from a symbolic representation to a metarepresentation exemplifies the process that Turchin has called a *metasystem transition* (Turchin, 1977, 1995, Heylighen, 1991, 1995). As we saw, this is the evolutionary emergence of a system that is capable of manipulating or controlling the organization of the system at the level below, thus extending its capabilities. The origin of human language and thus of RSS was such a metasystem transition, as it gave us the ability to manipulate some of the meanings that our neural networks generated. But it did not give us the ability to control the processes that generated such meaning: meanings were standardized through a social construction that produced a conventional RSS with its limited perspective on reality. Science, art, philosophy and spirituality have all attempted to extend our control over such sense-making processes, with concrete, albeit limited successes. The fundamental reason for the limitations is that sense-making processes are much more complex and fluid than expected, requiring an immense, ever-adapting network of interactions that extends across individual brains, bodies and world.

But this is where the on-going ICT revolution comes to the rescue. Present-day computers, memories, sensors and networks can handle immense flows of data, processes, and communications. They readily represent and thus externalize ideas, knowledge, images, sounds, movements, three-dimensional objects and most other phenomena that enter our consciousness. They can be programmed and manipulated, thus giving us control over these representations and processes. They have incorporated the sophisticated formal and operational modeling and data analysis methods of science, together with the powerful multimedia evocations of art; they have expanded our consciousness via prosthetic sensors, virtual and augmented reality; all this while being informed by the critical reflections of philosophy as to the nature of categories, cognition and representation. The only thing still lacking to complete the metasystem transition is *coordination*: these different tools and methods have largely developed independently, addressing different aspects and problems, in an apparently chaotic explosion of competing yet interconnected applications.

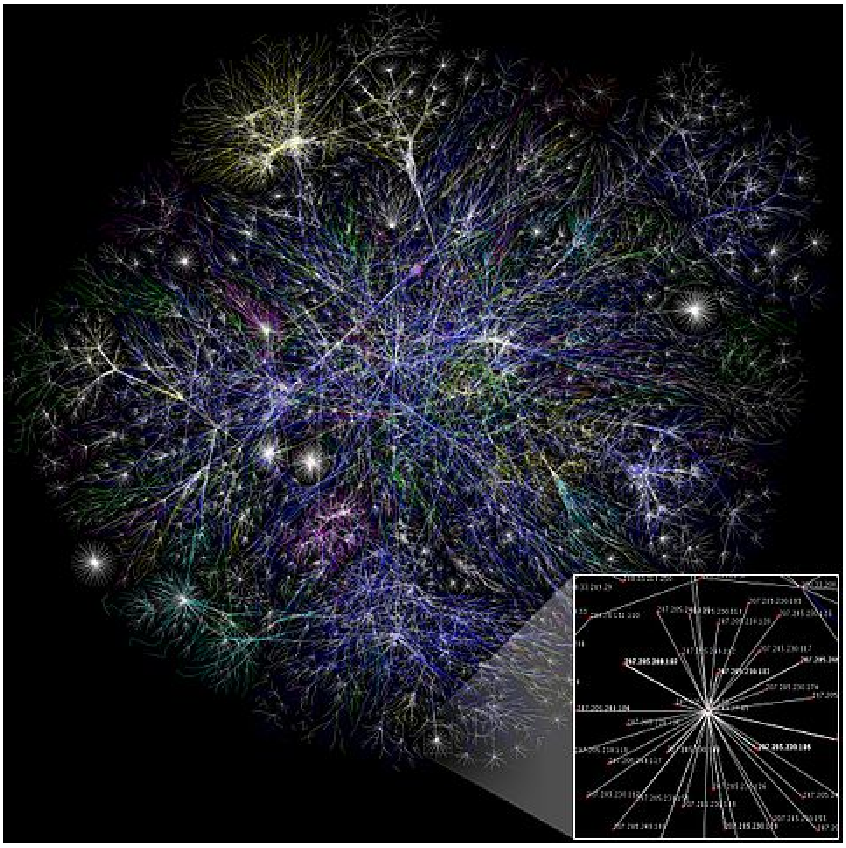

Fig. 9: Opte Project visualization of routing paths through a portion of the Internet. The network of connections could be seen as the pathways formed by neurons and synapses in a Global Brain that interconnects all humans and machines on the planet.

The "Global Brain" is the idea that a grand synthesis between all these different ICT applications is evolving, so as to produce the equivalent of an integrated neural network for humankind (Goertzel 2002; Heylighen 2011; Mayer-Kress & Barczys 1995) (see Fig. 9). The dynamics of this evolution is similar to the one that produced human language: interaction between incompatible applications initially produces misunderstanding, friction, and thus loss of information. This motivates the agents responsible for these applications to align their functions and standards with those of others, so that the systems become interoperable. Those that fail to do so will become increasingly irrelevant within the larger scheme, eventually losing the competition with those that use more compatible standards. This process of mutual alignment has been taking place since the beginnings of ICT. The Internet (TCP/IP) and World-Wide Web (HTML) standards are some of its great successes (Heylighen, 2017b). However, given the variety of novel applications that are constantly appearing and the complexity and unpredictability of the overall evolution, it is likely to take several more decades before a fully integrated system will take shape.

We need more than technological innovation and standardization to achieve such a grand synthesis of RSS-transcending approaches, though. We also need a deep theoretical or philosophical reflection that would lay the foundations for a different understanding of reality—one that is not biased by the RSS with its predisposition to reduce complex phenomena to static objects and categories that obey predetermined rules. This requires a completely new kind of ontology, one whose building blocks are not independent elements, but interconnected processes, actions or interactions (Rescher, 1996; Heylighen, 2011b). Here, meaning would not be a priori given like in an RSS, but self-organizing or emergent, so that new meanings can be generated without limit. Such *process ontology* may offer a new universal language that would unify all scientific and cultural disciplines, including art, philosophy and spirituality (Heylighen, 1984; Heylighen, Beigi and Veloz, 2015). This language would also become the standard medium used in ICT for the exchange and processing of knowledge and information across the most diverse types of hardware and software (Heylighen, 2017b). This function is already to some degree presaged by the *Semantic Web* standard (Berners-Lee and Fischetti, 1999; Allemang and Hendler, 2011), which functions as a cognitive extension of the present World-Wide Web, and a potential precursor of a Global Brain.

Let us try to conceive some of the likely attributes of the envisaged Global Brain. This superhuman intelligence will be fully collective or distributed (Heylighen, 2017a): no individual, computer or organization will have control over it. As the Internet makes it ever easier for information to flow from one person or machine to another one, that information will not stay in place but propagate across the network, just like activation spreads across the neurons of the human brain. Every source of information or processing, whether human or technological, can potentially contribute to any decision that needs to be made. Therefore, it is advantageous for any system to spread its net as widely as possible, incorporating as many information-producing agents as possible, without any one having a monopoly.

A derived attribute is what we might call "practical omniscience" (Heylighen, 2015): any information, knowledge or sensor that is available will eventually be integrated into the global network. More importantly, any knowledge that is not yet available but needed to solve some problem will, insofar possible, be generated on the spot, using the mechanisms of machine learning, model-building and creativity stimulation that AI, science and art have been developing. We saw that these mechanisms function by externalizing the sense-making processes of the human brain. But because the scope of the information they use will be planetary, they will far surpass the limits of the human brain. This is not just because of the ever-growing capabilities of ICT hardware, but because the distributed intelligence of the network will mobilize the collective perception, intuition and creativity of billions of human beings. As increasingly educated people with a great diversity of different experiences get access to all the knowledge, creative tools, metarepresentations, and supports for expanded consciousness, their capabilities for coming up with new concepts and insights, individually and collectively, are likely to explode. Through the planetary communication network, any new idea, generated by anyone or anything, can be immediately combined with any other idea or information to generate ever further insights, in an endless cascade of creative discovery.

This creative process would be similar to the scientific, technical and cultural progress that was enabled by human language, which we may see as the first level of symbolic representation. However, thanks to ICT and the other innovations we sketched, this process will have accelerated to such a degree that its results will seem nearly instantaneous. Extrapolations of accelerating technological advances indeed suggest that in the foreseeable future we would reach an apparently infinite speed of innovation—an event that has been dubbed the "technological Singularity" (Kurzweil, 2005; Goertzel and Goertzel, 2015; Heylighen, 2015). In mathematics, a singularity is a point where a continuous function becomes discontinuous, implying that we cannot extrapolate the curve beyond this singular point. But that also means that it becomes nearly impossible to infer what the consequences of this metasystem transition towards a supra-human intelligence will be…

We have come to the point where our description of the ongoing evolutionary transition towards a higher level of symbolization must stop. While I have tried to survey some of the social, technological and economic implications of this transition elsewhere (Heylighen, 2007, 2015), our focus here is on communication, cognition, and ultimately consciousness. But the complexity and change we are experiencing in this early 21st century is so overwhelming that no scientific model, artistic depiction, philosophical analysis, or personal awareness is as yet able to capture their full effect. We can only hope that the processes I have sketched will indeed self-organize into an integrated system of distributed intelligence, a Global Brain, that will give humankind a vastly broader and deeper insight into reality than it hitherto managed to achieve and thus help it to cope with the immense challenges that our planet has to deal with.

Glossary

symbol: a token, such as a word, sound, or graphic, that conventionally stands for some external phenomenon: the symbol's reference or meaning.

Rational Symbol System (RSS): a finite collection of symbols complemented by rules for combining those symbols into an infinite number of expressions that can represent meaningful situations. A typical example of an RSS is a verbal language with its words (symbols), grammar and logic (rules), and sentences (expressions). An RSS allows us to register, communicate, and reason about situations even when they are purely hypothetical. RSS use typifies human, rational intelligence as distinguished from animal intelligence.

Information and Communication Technologies (ICT): the whole of computers, memories, communication links, sensors and software that allow us to store, transfer, share and process symbolically expressed information, and thus automate the functions of an RSS.

distributed intelligence or **distributed cognition**: information-processing or problem-solving that is not localized in a particular component, such as an individual neuron, brain or computer, but that is spread out across a vast array of communicating components that work in parallel.

physical symbol system hypothesis: the founding assumption of "symbolic" theories of cognition and Artificial Intelligence (AI). It states that an RSS is all that is needed to support intelligence, whether in the human brain or in a computer.

symbol grounding problem: a fundamental shortcoming of symbolic theories of cognition. It notes that manipulating symbols according to abstract rules cannot tell us how these symbols are "grounded" in the real world, i.e. what they mean or which external phenomena they represent.

situated and embodied cognition: a complementary approach to symbolic theories of cognition that addresses the symbol grounding problem. Its guiding principle is that an intelligent agent needs a physical body to interact with the outside world, and not just an RSS. Meaning is grounded in the feedback loop that runs from the agent's perceptions of the world via its sensors to its actions that affect this world via its effectors, and back to its perception of these effects.

neural network: an abstract model of the way our brain processes information: activation from sensors propagates along complex networks of neurons and their connecting synapses, while changing shape, until it activates effectors. Neural networks learn new associations, patterns and knowledge by reinforcing successful connections and weakening unsuccessful ones. Thus, knowledge, processing and meaning are distributed across the whole network of neurons rather than localized in discrete symbols or rules.

subsymbolic cognition: the intuitive, fluid and holistic manner in which neural networks process information and make inferences, as contrasted with the rigid, discrete, logical procedures by which symbols are processed by a RSS.

ontology: the system of fundamental categories or distinctions between external phenomena that is implied by a particular RSS. These categories and their relations function as building blocks for our understanding of reality.

social construction of reality: the creation of an RSS and its associated ontology through social interactions between agents. With each interaction, symbols and categories shared between the agents are reinforced, while dissimilar ones are made more similar or eliminated, until all agents agree on all symbols, categories and rules. The resulting system is rigid, limited and to some degree arbitrary, as it depends on the local group or culture.

formalization: the method used in science to explicitly specify the meaning of a symbol in terms of other symbols, ideally by means of logical, computational or mathematical definitions and axioms.

operationalization: the method used in science to explicitly specify the meaning of a symbol in terms of the physical operations, such as measurements, tests or experiments, that can be used to establish in how far a phenomenon belongs to the category denoted by the symbol.

limitation principles: a variety of results achieved in 20th century science, including the theorem of Gödel, the Heisenberg uncertainty principle, and the halting problem, that demonstrate the impossibility of reaching a complete representation of reality through either formalization or operationalization.

linguistic complementarity principle: a generalization of several of these limitation principles as well as the symbol grounding problem. It notes that no language or RSS can fully describe its own interpretation, i.e. the processes by which its symbols are to be connected to the external phenomena they represent.

icon: a symbol, such as portrait, picture or diagram, that conveys its meaning by its perceived resemblance to the phenomenon it depicts.

metaphor: a symbol that conveys its meaning indirectly, through the analogy that exists between the phenomenon it normally refers to, and the one it metaphorically refers to.

surrealism: a movement in art that questions our standard representation of reality by exploring "impossible" or "irrational" combinations of symbols, thus evoking the subsymbolic associations typical of dreams and the subconscious.

spirituality or **mysticism**: the quest for a more direct access to the fluid experience of an undivided reality that is characteristic of subsymbolic cognition. This is normally suppressed or obscured by the rigid categorizations of the RSS. Typical methods proposed to achieve such "expanded consciousness" or "enlightenment" are meditation, mindfulness and psychedelic drugs.

virtual reality: visual, auditory and sometimes tactile simulation of a three-dimensional world by means of a computer application. The simulated perception reacts to the movements of the user as if the user were moving in a real environment.

augmented reality: computer-generated information about the real world that is superposed on a user's actual perception of that world .

metasystem transition: the evolutionary emergence of a new level of organization or intelligence in a system. After the transition, the fixed rules that hitherto governed the system's activity have become themselves subjected to adaptation and control by higher-order rules. The emergence of RSS-governed symbolic cognition out of subsymbolic cognition is an example of a metasystem transition.

transcending the RSS: the metasystem transition that would overcome most of the remaining limitations of symbolic cognition, by allowing us to automatically adapt and create symbolic categories and rules depending on the needs of the situation. Science, art, philosophy, spirituality and their extensions via ICT can all be seen as contributing to the creation of such a "superhuman" level of cognition.

Global Brain: a metaphor for the putative suprahuman, distributed intelligence that would emerge by the integration of all human and machine intelligences as connected by the planetary ICT network. The Global Brain would form the substrate for the RSS-transcending, next level of cognition, by providing the combined knowledge, creativity, intuition, memory and processing capability necessary to on the spot generate the concepts or categories, rules, and models that would be needed to solve any problem.

References

Adamowicz, E. (1998) *Surrealist Collage in Text and Image: Dissecting the Exquisite Corpse*. Cambridge University Press.

Ashby, W. R. (1964) *An Introduction to Cybernetics*. Methuen London.

Axelrod, R. (1997) 'The dissemination of culture', *Journal of conflict resolution*, 41(2), p. 203.

Bailer-Jones, D. M. (2009) *Scientific Models in Philosophy of Science*. University of Pittsburgh Pre.

Barrow, J. D. (1998) *Impossibility: The limits of science and the science of limits*. Oxford University Press, USA.

Bechtel, W. and Abrahamsen, A. (1991) *Connectionism and the mind: An introduction to parallel processing in networks.* Basil Blackwell.

Bechtel, W. and Graham, G. (eds) (1999) 'A Companion to Cognitive Science', in *A Companion to Cognitive Science*. Oxford, UK: Blackwell Publishing Ltd, pp. 506–517. doi: 10.1111/b.9780631218517.1999.00042.x.

Berger Peter, L. and Luckmann, T. (1966) 'The social construction of reality: A treatise in the sociology of knowledge', *Garden City, NY: First Anchor*.

Bickhard, M. H. and Terveen, L. (1995) *Foundational issues in artificial intelligence and cognitive science: Impasse and solution*. North Holland.

Bishop, S. R. *et al.* (2004) 'Mindfulness: A Proposed Operational Definition', *Clinical Psychology: Science and Practice*, 11(3), pp. 230–241. doi: 10.1093/clipsy.bph077.

Black, M. (1979) 'Wittgenstein's Language-games', *Dialectica*, pp. 337–353.

Bodnar, I. (2018) 'Aristotle's Natural Philosophy', in Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy*. Spring 2018. Metaphysics Research Lab, Stanford University. Available at: https://plato.stanford.edu/archives/spr2018/entries/aristotle-natphil/ (Accessed: 12 July 2018).

Bohil, C. J., Alicea, B. and Biocca, F. A. (2011) 'Virtual reality in neuroscience research and therapy', *Nature Reviews Neuroscience*, 12(12), p. 752. doi: 10.1038/nrn3122.

Boulding, K. E. (1956) 'General Systems Theory-The Skeleton of Science', *Management Science*, 2(3), pp. 197–208.

Burkholder, L. (1987) 'The Halting Problem', *SIGACT News*, 18(3), pp. 48–60. doi: 10.1145/24658.24665.

Burks, A. W. (1949) 'Icon, index, and symbol', *Philosophy and phenomenological research*, 9(4), pp. 673–689.

Chandler, D. (2017) *Semiotics: The Basics*. Routledge.

Clark, A. (2008) *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford University Press.

Copeland, B. J. (2002) 'The Genesis of Possible Worlds Semantics', *Journal of Philosophical Logic*, 31(2), pp. 99–137. doi: 10.1023/A:1015273407895.

Deikman, A. J. (1966) 'De-automatization and the Mystic Experience.', *Psychiatry: Journal for the Study of Interpersonal Processes*. Available at: http://psycnet.apa.org/psycinfo/1967-16094-001 (Accessed: 8 October 2014).

Deikman, A. J. (2000) 'A functional approach to mysticism', *Journal of Consciousness Studies*, 7(11/12), pp. 75–92.

Dreyfus, H. L. (1992) *What computers still can't do: a critique of artificial reason*. The MIT Press.

Edwards, D. (2008) *Artscience: Creativity in the post-Google generation*. Harvard University Press.

Fay, N. *et al.* (2010) 'The Interactive Evolution of Human Communication Systems', *Cognitive Science*, 34(3), pp. 351–386. doi: 10.1111/j.1551-6709.2009.01090.x.

Ferber, J. (1998) *Multi-agent systems: an introduction to distributed artificial intelligence*. Addison-Wesley.

Folina, J. (2010) 'Poincaré's Philosophy of Mathematics', *Internet Encyclopedia of Philosophy*. Available at: http://www.iep.utm.edu/poi-math/ (Accessed: 12 March 2012).

Franzén, T. (2005) *Gödel's theorem: an incomplete guide to its use and abuse*. AK Peters Wellesley, MA.

Fusaroli, R. and Tylén, K. (2012) 'Carving language for social coordination: A dynamical approach', *Interaction Studies*, 13(1), pp. 103–124. doi: 10.1075/is.13.1.07fus.

Garrod, S. and Doherty, G. (1994) 'Conversation, co-ordination and convention: An empirical investigation of how groups establish linguistic conventions', *Cognition*, 53(3), pp. 181–215.

Glenn, J. C. (1989) *Future Mind: Artificial Intelligence: The Merging of the Mystical and the Technological in the 21st Century*. Acropolis Books, Incorporated.

Glock, H.-J. (2008) *What is Analytic Philosophy?* Cambridge University Press.

Gödel, K. (1992) *On formally undecidable propositions of Principia Mathematica and related systems*. Courier Corporation.

Goertzel, B. (2002) *Creating internet intelligence: Wild computing, distributed digital consciousness, and the emerging global brain*. Kluwer Academic/Plenum Publishers.

Goertzel, B. and Goertzel, T. (2015) *The End of the Beginning: Life, Society and Economy on the Brink of the Singularity*. 1 edition. Humanity+ Press.

Goertzel, B. and Pennachin, C. (2007) *Artificial general intelligence*. Springer Verlag.

Harnad, S. (1990) 'The symbol grounding problem', *Physica D: Nonlinear Phenomena*, 42(1–3), pp. 335–346.

Harnad, S. (2002) 'Symbol grounding and the origin of language', *Computationalism: new directions*, pp. 143–158.

Harris, S. (2014) *Waking Up: A Guide to Spirituality Without Religion*. Simon and Schuster.

Hawkes, T. (2003) *Structuralism and semiotics*. Routledge..

He, B. *et al.* (2013) 'Brain–Computer Interfaces', in He, B. (ed.) *Neural Engineering*. Boston, MA: Springer US, pp. 87–151..

Henrich, J. and Boyd, R. (1998) 'The Evolution of Conformist Transmission and the Emergence of Between-Group Differences', *Evolution and Human Behavior*, 19(4), pp. 215–241.

Heylighen, F. (1984) 'Transcending the Conceptual-Symbolic Code', *'O': tijdschrift voor filosofische beproevingen*, 7. Available at: http://pespmc1.vub.ac.be/papers/TranscendingCS-code.pdf (Accessed: 22 April 2016).

Heylighen, F. (1988) 'Formulating the problem of problem-formulation', in *Cybernetics and Systems '88*. Vienna: Austrian Society For Cybernetics, pp. 949–957. Available at: http://pespmc1.vub.ac.be/Papers/Formulating.pdf.

Heylighen, F. (1990) *Representation and Change: A Metarepresentational Framework for the Foundations of Physical and Cognitive Science*. Communication and Cognition. Available at: http://pespmc1.vub.ac.be/books/Rep%26Change.pdf.

Heylighen, F. (1991) 'Cognitive Levels of Evolution: from pre-rational to meta-rational', in Geyer, F. (ed.) *The Cybernetics of Complex Systems-Self-organization, Evolution and Social Change*, pp. 75–91. Available at: http://pespmc1.vub.ac.be/Papers/CognLevels.pdf.

Heylighen, F. (1995) '(Meta)systems as constraints on variation— a classification and natural history of metasystem transitions', *World Futures*, 45(1), pp. 59–85. doi: 10.1080/02604027.1995.9972554.

Heylighen, F. (1999) 'Advantages and limitations of formal expression', *Foundations of Science*, 4(1), pp. 25–56. doi: 10.1023/A:1009686703349.

Heylighen, F. (2001) 'Mining Associative Meanings from the Web: from word disambiguation to the global brain', in Lutjeharms, M. and Temmerman, R. (eds) *Proceedings of Trends in Special Language and Language Technology*. De Boeck Hoger, pp. 15–44. Available at: http://pespmc1.vub.ac.be/papers/MiningMeaning.pdf.

Heylighen, F. (2007) 'The Global Superorganism: an evolutionary-cybernetic model of the emerging network society', *Social Evolution & History*, 6(1), pp. 58–119.

Heylighen, F. (2011) 'Conceptions of a Global Brain: an historical review', in *Evolution: Cosmic, Biological, and Social, eds. Grinin, L. E., Carneiro, R. L., Korotayev A. V., Spier F.* Uchitel Publishing, pp. 274–289. Available at: http://pcp.vub.ac.be/papers/GB-conceptions-Rodrigue.pdf.

Heylighen, F. (2013) 'Self-organization in Communicating Groups: the emergence of coordination, shared references and collective intelligence', in Massip-Bonet, À. and Bastardas-Boada, A. (eds) *Complexity Perspectives on Language, Communication and Society*. Berlin, Germany: Springer (Understanding Complex Systems), pp. 117–149. Available at: http://pcp.vub.ac.be/Papers/Barcelona-LanguageSO.pdf (Accessed: 19 November 2012).

Heylighen, F. (2015) 'Return to Eden? Promises and Perils on the Road to a Global Superintelligence', in Goertzel, B. and Goertzel, T. (eds) *The End of the Beginning: Life, Society and Economy on the Brink of the Singularity*. Humanity+ Press. Available at: http://pespmc1.vub.ac.be/Papers/BrinkofSingularity.pdf.

Heylighen, F. (2017a) 'Distributed Intelligence Technologies: present and future applications', in Burgin, M. and Hofkirchner, W. (eds) *The Future Information Society: Social and Technological Problems*. World Scientific, pp. 179–212. Available at: http://pespmc1.vub.ac.be/Papers/GB-applications-survey-WS.pdf.

Heylighen, F. (2017b) 'The Offer Network Protocol: mathematical foundations and a roadmap for the development of a global brain', *The European Physical Journal Special Topics*, 226(2), pp. 283–312. doi: 10.1140/epjst/e2016-60241-5.

Heylighen, F., Beigi, S. and Veloz, T. (2015) *Chemical Organization Theory as a modeling framework for self-organization, autopoiesis and resilience*. ECCO Working Papers 2015-01. Available at: http://pespmc1.vub.ac.be/Papers/COT-applicationsurvey.pdf.

Heylighen, F., Cilliers, P. and Gershenson, C. (2007) 'Complexity and Philosophy', in Bogg, J. and Geyer, R. (eds) *Complexity, science and society*. Oxford: Radcliffe Publishing, pp. 117–134. Available at: http://arxiv.org/abs/cs/0604072.

Heylighen, F. and Joslyn, C. (2003) 'Cybernetics and Second-Order Cybernetics', in Robert A. Meyers (ed.) *Encyclopedia of Physical Science and Technology*. 3rd edn. New York: Academic Press, pp. 155–169. Available at: http://pespmc1.vub.ac.be/Papers/Cybernetics-EPST.pdf (Accessed: 28 November 2012).

Hilborn, R. C. (2004) 'Sea gulls, butterflies, and grasshoppers: A brief history of the butterfly effect in nonlinear dynamics', *American Journal of Physics*, 72(4), pp. 425–427.

Hilgevoord, J. and Uffink, J. (2001) 'The Uncertainty Principle'. Available at: https://stanford.library.sydney.edu.au/entries/qt-uncertainty/ (Accessed: 24 August 2018).

Huxley, A. (2009) *The Doors of Perception and Heaven and Hell*. Later Printing edition. New York: Harper Perennial Modern Classics.

Kelley, T. D. (2003) 'Symbolic and Sub-Symbolic Representations in Computational Models of Human Cognition: What Can be Learned from Biology?, *Theory & Psychology*, 13(6), pp. 847–860. doi: 10.1177/0959354303136005.

Kinny, D. and Georgeff, M. (1997) 'Modelling and design of multi-agent systems', *Lecture Notes In Computer Science*, pp. 1–20.

Kohn, L. (2008) *Meditation works: In the daoist, Buddhist and Hindu traditions*. Three Pines Press.

Korzybski, A. (1958) *Science and Sanity: An Introduction to Non-Aristotelian Systems and General Semantics*. Institute of GS.

Krueger, M., Packer, R. and Jordan, K. (2001) *Multimedia: From Wagner to Virtual Reality*. 1st ed. New York: Norton.

Kupers, R. and Ptito, M. (2004) '"Seeing" through the tongue: cross-modal plasticity in the congenitally blind', *International Congress Series*. (Frontiers in Human Brain Topology. Proceedings of ISBET 2004), 1270, pp. 79–84. doi: 10.1016/j.ics.2004.04.053.

Kurzweil, R. (2005) *The singularity is near*. Penguin books.

Löfgren, L. (1992) 'Complementarity in Language; Toward a General Understanding', in *Nature, Cognition and System II*. Springer, Dordrecht (Theory and Decision Library), pp. 113–153. doi: 10.1007/978-94-011-2779-0_8.

Ludwig, D. S. and Kabat-Zinn, J. (2008) 'Mindfulness in Medicine', *JAMA*, 300(11), pp. 1350–1352. doi: 10.1001/jama.300.11.1350.

Massé, A. B. *et al.* (2008) 'How is meaning grounded in dictionary definitions?', in *Proceedings of the 3rd Textgraphs Workshop on Graph-Based Algorithms for Natural Language Processing*. Association for Computational Linguistics, pp. 17–24.

Matthews, J. H. (1986) *Languages of Surrealism*. University of Missouri Press.

Mayer-Kress, G. and Barczys, C. (1995) 'The global brain as an emergent structure from the Worldwide Computing Network, and its implications for modeling', *The information society*, 11(1), pp. 1–27.

Maynard Smith, J. and Szathmáry, E. (1997) *The major transitions in evolution*. Oxford University Press, USA.

McLeod, P., Plunkett, K. and Rolls, E. T. (1998) *Introduction to connectionist modelling of cognitive processes*. Oxford University Press Oxford.

Merrell, F. (2001) 'Charles Sanders Peirce's concept of the sign', *The Routledge companion to semiotics and linguistics*, pp. 28–39.

Mesarović, M. D. and Takahara, Y. (1975) *General systems theory: mathematical foundations*. Elsevier Science.

Nauta, D. (1972) *The Meaning of Information*. The Hague: Mouton.

Newell, A. and Simon, H. A. (1972) *Human problem solving*. Prentice-Hall Englewood Cliffs, NJ.

Newell, A. and Simon, H. A. (1976) 'Computer science as empirical inquiry: Symbols and search', *Communications of the ACM*, 19(3), pp. 113–126.

Nilsson, N. J. (2007) 'The Physical Symbol System Hypothesis: Status and Prospects', in *50 Years of Artificial Intelligence*. Springer, Berlin, Heidelberg (Lecture Notes in Computer Science), pp. 9–17. doi: 10.1007/978-3-540-77296-5_2.

Noë, A. (2010) *Out of Our Heads: Why You Are Not Your Brain, and Other Lessons from the Biology of Consciousness*. Farrar, Straus and Giroux.

Parsons, W. B. (1999) *The Enigma of the Oceanic Feeling: Revisioning the Psychoanalytic Theory of Mysticism*. Oxford University Press.

Petrovic, K. (2018) *On Cosmogony*. ECCO Working Papers 2018-02. Available at: http://katarinapetrovic.net/on-cosmogony.pdf.

Poggioli, R. (1981) *The Theory of the Avant-Garde*. Translated by G. Fitzgerald. Cambridge, Mass.: The Belknap Press of Harvard University Press.

Root-Bernstein, B. *et al.* (2011) 'ArtScience: integrative collaboration to create a sustainable future', *Leonardo*, 44,(3), p. 192.

Russell, S. J. and Norvig, P. (2009) *Artificial intelligence: a modern approach*. Prentice hall.

Sanchez-Vives, M. V. and Slater, M. (2005) 'Opinion: From presence to consciousness through virtual reality', *Nature Reviews Neuroscience*, 6(4), pp. 332–339. doi: 10.1038/nrn1651.

Schartner, M. M. *et al.* (2017) 'Increased spontaneous MEG signal diversity for psychoactive doses of ketamine, LSD and psilocybin', *Scientific Reports*, 7. doi: 10.1038/srep46421.

Schmidhuber, J. (2015) 'Deep learning in neural networks: An overview', *Neural Networks*, 61, pp. 85–117. doi: 10.1016/j.neunet.2014.09.003.

Sessa, D. B. (2013) *The Psychedelic Renaissance: Reassessing the Role of Psychedelic Drugs in 21st Century Psychiatry and Society*. 1st edition. London: Muswell Hill Press.

Shannon, C. E. and Weaver, W. (1963) *The Mathematical Theory of Communication*. University of Illinois Press.

Siler, T. (2011) 'The artscience program for realizing human potential', *Leonardo*, 44(5), pp. 417–424.

Smith, H. and Marranca, R. (2009) *The world's religions*. HarperOne New York.

Steels, L. and Belpaeme, T. (2005) 'Coordinating perceptually grounded categories through language: A case study for colour', *Behavioral and brain sciences*, 28(04), pp. 469–489.

Stutley, M. (2003) *Shamanism: An Introduction*. Routledge. doi: 10.4324/9780203398166.

Thagard, P. (1996) *Mind: Introduction to cognitive science*. MIT press Cambridge, MA.

Turchin, V. (1977) 'The Phenomenon of Science. A Cybernetic Approach to Human Evolution', *New York: Columbia University*.

Turchin, V. (1995) 'A dialogue on metasystem transition', *World Futures*, 45, pp. 5–57.

Van Krevelen, D. W. F. and Poelman, R. (2010) 'A survey of augmented reality technologies, applications and limitations', *International Journal of Virtual Reality*, 9(2), p. 1.

West, D. (2000) *An Introduction to Continental Philosophy*. Cambridge England ; Cambridge, MA: Polity Press.

Wiener, N. (1965) *Cybernetics: or, Control and Communication in the Animal and the Machine*. The MIT Press.

Wittgenstein, L. (2010) *Philosophical investigations*. John Wiley & Sons.

Wood, P. (2004) *Art of the Avant-gardes*. New Haven, Conn.: Yale University Press.

Yanofsky, N. S. (2013) *The outer limits of reason: what science, mathematics, and logic cannot tell us*. MIT Press.