

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/301803762>

Non-dualism

Data · May 2016

CITATIONS

0

READS

240

2 authors:



Francis Heylighen

Vrije Universiteit Brussel

157 PUBLICATIONS 5,457 CITATIONS

[SEE PROFILE](#)



Shima Beigi

University of Oxford

24 PUBLICATIONS 31 CITATIONS

[SEE PROFILE](#)

Some of the authors of this publication are also working on these related projects:



<http://pespmc1.vub.ac.be/ECCO/ECCO-Papers/GB-Protocol-proposal.pdf> [View project](#)



Mindfulness Engineering in Action [View project](#)

Mind outside Brain: a radically non-dualist foundation for distributed cognition

Francis Heylighen & Shima Beigi

Global Brain Institute

Vrije Universiteit Brussel

fheyligh@vub.ac.be, shima.beigi@gmail.com

Abstract: We approach the problem of the extended mind from a radically non-dualist perspective. The separation between mind and matter is an artefact of the outdated mechanistic worldview, which leaves no room for mental phenomena such as agency, intentionality, or feeling. We propose to replace it by an action ontology, which conceives mind and matter as aspects of the same network of processes. By adopting the intentional stance, we interpret the catalysts of elementary reactions as agents exhibiting desires, intentions, and sensations. Autopoietic networks of reactions constitute more complex super-agents, which moreover exhibit memory, deliberation and sense-making. In the specific case of social networks, individual agents coordinate their actions via the propagation of challenges. The distributed cognition that emerges from this interaction cannot be situated in any individual brain. This non-dualist, holistic view extends and operationalizes process metaphysics and Eastern philosophies. It is supported by both mindfulness experiences and mathematical models of action, self-organization, and cognition.

Introduction

This paper is part of a volume focused on socially extended knowledge. Extended knowledge is an aspect of the more general thesis of the *extended mind* (Clark & Chalmers, 1998), which states that mental phenomena, such as memory, knowledge and sensation, extend outside the individual human brain, and into the material and social environment. In other words, the skull can no longer be seen as a clear physical boundary between (inside) mind and (outside) world.

While the extended mind hypothesis originates in philosophy, a number of closely related conceptions have been formulated in cognitive science under headers such as *situated*, *embodied*, *enactive* and *embedded cognition* (Anderson, 2003; Clark, 1998; Stewart, Gapenne, & Paolo, 2014; Susi & Ziemke, 2001). The general idea is that human cognition is not confined to information processing within the brain, but actively dependent on external phenomena. These include the body, cognitive tools such as notebooks and computers, the situation, the interactions between agent and environment, communications with other agents, and social systems. We will summarize this broad scale of “extensions” under the header of

distributed cognition (Hutchins, 2000), as they all imply that cognitive content and processes are distributed across a variety of agents, objects and actions. Only some of those are located inside the human brain; yet all of them contribute to human decisions by providing part of the information necessary to make these decisions.

While the distributed nature of information processing is difficult to deny, the extended mind thesis remains controversial. The reason seems to be that most philosophers investigating this idea feel that there is a fundamental difference between truly “mental” phenomena, such as belief, desire or intention, and the merely “mechanical” phenomena of information transmission, storage and processing. Thus, the Alzheimer patient Otto, who relies on his notebook as an external memory in the original “extended mind” thought experiment (Clark & Chalmers, 1998), does not really seem to outsource his desires, intentions or beliefs to his notebook. He merely believes (internally) that this notebook is a dependable tool for storing information (externally) that his own brain cannot reliably store. These and other intuitions about how the mind works fuel an on-going discussion about whether and in how far mental phenomena truly can extend outside of the brain.

The aim of the present paper is to propose a radical resolution to this controversy: we assume that *mind* is a ubiquitous property of all minimally active matter (Heylighen, 2011). It is in no way restricted to the human brain—although that is the place where we know it in its most concentrated form. Therefore, the extended mind hypothesis is in fact misguided, because it assumes that the mind originates in the brain, and merely “extends” itself a little bit outside in order to increase its reach, the way one’s arm extends itself by grasping a stick. While ancient mystical traditions and idealist philosophies have formulated similar panpsychist ideas (Seager, 2006), the approach we propose is rooted in contemporary science—in particular cybernetics, cognitive science, and complex systems theory. As such, it strives to formulate its assumptions as precisely and concretely as possible, if possible in a mathematical or computational form (Heylighen, Busseniers, Veitas, Vidal, & Weinbaum, 2012), so that they can be tested and applied in real-world situations—and not just in the thought experiments beloved by philosophers.

But before we can elaborate our thesis of a ubiquitously distributed mind, we need to explain why this idea appears so radical, and why the comparatively modest hypothesis of extended mind or extended knowledge remains so controversial. For this we need to go back to what we see as the root of the problem: Cartesian dualism and Newtonian mechanics.

From dualism to action ontology

Descartes formulated his philosophy of the fundamental duality of mind and matter in the context of the mechanistic worldview that was emerging at the time. In Descartes’s view, the body was merely a complicated mechanical system, an automaton essentially similar to a clockwork in which one gear passes on its movement to another gear. This understanding was

later elaborated scientifically by Newton, Laplace and their successors as the foundation for the mechanistic worldview (Toulmin, 1993) that became dominant in the 19th century. We will from now on refer to this mechanistic vision of the world as the *Newtonian worldview* (Heylighen, Cilliers, & Gershenson, 2007). By investigating its assumptions, we will try to clarify why Descartes and many thinkers after him felt they had to introduce mind as a realm separate from the realm of matter.

The Newtonian worldview reduces the world to a collection of material objects that move through space along fixed trajectories. The laws of mechanics specify how one material body exerts a force on another material body, thus affecting its movement—similarly to the way a gear transmits motion to another gear. These laws completely determine the trajectories of all material bodies—just like the movement of a clockwork is rigidly fixed by the configuration of its gears. Therefore, Laplace’s imaginary demon, who is able to precisely observe all the positions and velocities of all the pieces of matter in the universe, would be able to predict all their future movements, and thus all change or evolution. Such determinism implies that there is no freedom to intervene, to act intentionally, or to choose between different courses of action.

For Descartes, it was obvious that the mind has such freedom. Therefore, the mind cannot be subjected to mechanical laws. But since all matter obeys such laws, the mind cannot be material. Therefore, the mind must be independent, belonging to a realm separate from the realm of matter. This in principle allows the mind to leave its material body—the way the soul is supposed to do in the religious conception of dying. However, this assumption immediately creates a paradox: if mind and matter are independent, then how can the one affect the other? The seemingly unsolvable mind-body problem (McGinn, 1989) is in essence a series of variations on the following questions: 1) how can the immaterial mind sense, feel or become conscious of the material world?; 2) how can this mind control the matter of the body, given that this matter is already fully controlled by mechanistic laws?

The strict Cartesian separation between mind and matter has been abandoned by all but a few contemporary scientists and philosophers. Most scholars nowadays agree that the mind supervenes on the matter of the brain, i.e. it cannot exist without the presence of this material infrastructure. Thus, few academics still believe in the existence of an immaterial soul. However, in practice most of them still stick to what Dennett has called “Cartesian materialism” (Dennett & Kinsbourne, 1997). This is the implicit assumption that while the mind is somehow constituted out of the matter in the brain, it still has some kind of autonomous agency that separates it from the rest of the world. This intuition is based on the same apparent inconsistency observed by Descartes between the mechanistic view of the world and our experience of free will—however without offering any resolution to the paradox.

This intuitive separation between mind and world is reinforced by what Chalmers has called the “hard problem of consciousness” (Chalmers, 1995). The mind does not only freely decide and act, it also subjectively experiences the world; it “feels” the phenomena it

encounters. If the mind were merely a mechanical system, then it seems that there would be no room for such subjective sensation or phenomenal consciousness. The only way to affect a mechanical system is to change the way its material components move; the matter itself however remains inert, lifeless, insensitive. Chalmers illustrates this problem with the zombie thought experiment—where a zombie can be seen as a robot-like creature in which incoming stimuli are transmitted and processed through the forces inside the mechanism, eventually producing outgoing actions. Thus, a zombie is merely a more sophisticated, intelligent version of the automaton conceived by Descartes as a model of the body. Assuming that we could build a zombie that is not distinguishable in its behavior from a real human person, then—the argument goes—that zombie would still be lacking something essential, namely phenomenal consciousness. Though it may move, act and react like a person, it cannot feel like a person. Thus, the mind somehow still has this mysterious property of feeling that is absent in matter.

Thus, we see that modern conceptions of mind are still implicitly dualist, even though few would deny its materialist basis. Our position is that this separation is an artefact of the Newtonian worldview. Mind seems incompatible with the material world because we have a much too simple, outdated view of what that world really consists of.

Modern physics, chemistry and biology have long abandoned the static and deterministic Newtonian worldview. Instead, they see the world as a network of processes. The “matter” that constitutes it is completely unlike the inert “billiard-ball”-like particles following predetermined trajectories envisaged by Newtonian mechanics. Instead, quantum field theories see particles as merely temporary, local “condensations” of fields representing potential interactions. Particles are constantly being created and destroyed by elementary particle reactions, and this in a fundamentally unpredictable manner. For example, it is intrinsically impossible to predict when a radioactive atom will decay into smaller particles. They can even emerge out of nothing as virtual particle-antiparticle pairs formed by quantum fluctuations of the vacuum.

At a higher level, the molecules that constitute living organisms are similarly ephemeral, constantly being produced and consumed by the chemical reactions that constitute the organism’s metabolism. Cells and organelles in the body too are in a constant flux, being broken down by processes such as apoptosis and autophagy, while new ones are grown through cell division and from stem cells. The same processes can again be found at the level of ecosystems, where relations of predation, symbiosis and reproduction between organisms and species join with meteorological and geological forces to produce a constantly changing landscape of resources and constraints, opportunities and dangers.

All these processes are indeterministic, both in principle—because of the underlying Heisenberg uncertainty principle—and in practice—because of the chaotic “butterfly effects” that magnify unobservably small perturbations into macroscopic and sometimes catastrophic changes (Heylighen, 2012; Heylighen et al., 2007). However, these processes are not random, but *evolutionary*: they have a preferred direction towards survival and growth (fitness), engendering increasingly complex and intelligent forms of organization (Heylighen, 1999).

Thus, they lead to the emergence of ever more sophisticated, meaningful and adaptive forms. This evolutionary worldview is very different from the lifeless, static picture of the clockwork universe, where inert pieces of matter follow predetermined trajectories. In such an evolving, interconnected world, mind no longer appears like an alien entity that cannot be explained by scientific principles, but rather as a natural emanation of the way networks of processes self-organize into goal-directed, adaptive systems.

This is of course not a novel idea. It has been formulated by philosophers such as Whitehead, Bergson and Teilhard de Chardin under the label of *process metaphysics* (Rescher, 1996; Teilhard de Chardin, 1959; Whitehead, 1978). But analytically trained philosophers are understandably not very keen on these rather mystical and obscure theories, preferring the clear distinctions of logic and mathematics to these poetic and grandiloquent writings. Therefore, analytic philosophy has tended to stay firmly rooted in the reductionist approach of Newtonian science. The problem is that this leads it straight back into an implicit dualism, and its apparently unsolvable mind-body problem.

The thesis of this paper is that you can have your cake and eat it: it is possible to develop an understanding of the mind that is both non-dual *and* analytic—in the sense of based on clearly defined, formal distinctions. To achieve that, we need to replace the vagueness of process metaphysics by the concreteness of what has been called *action ontology* (Heylighen, 2011; Turchin, 1993). That will allow us to “extend” the mind not just across notebooks and social systems, but across the whole of nature and society.

Agents and the intentional stance

Rather than starting from particles or pieces of matter, the action ontology is based on actions or reactions. These are elementary processes that lead from some initial condition X to a subsequent condition Y:

$$X \rightarrow Y$$

These conditions can in general be decomposed into conjunctions of more elementary conditions. Adopting the notation used for reactions in physics and chemistry, we will denote conjunctions by the “+” operator:

$$a + b + \dots \rightarrow e + f + \dots$$

A reaction can be interpreted in several, approximately equivalent manners: as a transition from the state X to the next state Y; as a causation producing the effect Y out of the cause X; as a production rule, to be read as “if X, then Y”, which specifies under which condition X the

action (change of condition) Y is produced; as a transformation of a situation characterized by X into a new situation characterized by Y.

Conditions merely specify that some distinguishable category of phenomena is present at the beginning or end of the reaction. Therefore, reactions can represent processes in any domain or discipline. This is best illustrated by a few examples, as listed in Table 1.

Elementary particle reaction	$n \rightarrow p + e^- + \nu_e$	(Beta decay of neutron)
Chemical reaction	$2H_2 + O_2 \rightarrow 2H_2O$	(production of water)
Ecological process	plants + sunlight + carbon dioxide + minerals \rightarrow more plants + oxygen	
Causal rule	Glass falls + hard floor \rightarrow Glass breaks	
Action of thermostat	Temperature $< 21^\circ$ \rightarrow switch on heating	
Animal action	dog + meat \rightarrow dog + meat eaten	
Human action	See friend \rightarrow greet friend	

Table 1: examples of reactions in different domains

In all these examples, the reaction starts from a distinguishable initial state, which is then transformed into a new state. While this may seem to make reactions dependent on states, in fact states can be defined in terms of the reactions that are possible in that state (Heylighen, 2011; Turchin, 1993). Thus, (re)actions or processes are truly the building blocks of the action ontology, while states are secondary.

Agents (A) can be defined in this framework as necessary conditions for the occurrence of a reaction, which however are not themselves affected by the reaction;

$$A + X \rightarrow A + Y$$

In chemistry, the function of A is the one of a catalyst: it enables the reaction that converts X into Y. Since A remains invariant during the reaction, but needs to be present in order for the reaction to take place, it can be seen as the agent of the conversion. The *reaction* between A, X and Y can therefore be reinterpreted as an *action* performed by the *agent* A on condition X in order to produce condition Y:

A: $X \rightarrow Y$

Agents will in general participate in different reactions. This means that they are able to perform different actions, reacting to different conditions by different actions producing different new conditions. For example:

A: $X \rightarrow Y, Y \rightarrow Z, U \rightarrow Z, V \rightarrow W, E \rightarrow F, F \rightarrow G, W \rightarrow G, \dots$

We are now ready to ascribe rudimentary mental properties to an agent. First, agents have “sensations”: they are able to sense the conditions to which they react, acting differently under different conditions (X, Y, U, ...). Inert pieces of matter do not react to specific conditions with specific actions: they are “insensitive” to their situation.

Second, agents have “desires” or “goals”. One way to understand this is by noting that the list of actions that an agent can perform defines a *dynamical system* (Beer, 1995, 2000). This is a mathematical model of a process that describes a trajectory in a state space, leading from some initial state (say X), to the next state (Y), and the next (Z), and so on. Dynamical systems typically have one or more *attractors*. These are states or subsets of states that are in a sense end points of the process: different trajectories lead into the attractor, but no trajectory leads out of it. In the example of agent A above, Z and G are attractors (see Fig. 1). Starting from X or Y, the system will end up in Z; starting from V, E, F, or W, the system will end up in G.




Figure 1: a phase portrait of the dynamical system defined by an agent. Arrows represent the agent’s actions leading from one state (e.g. X) to the next (e.g. Z). The shaded areas Z and G are attractors, each surrounded by their basin, from which all courses of action lead into the respective attractor. Curvy broken arrows represent external disturbances, which make the state deviate from its normal course of action

The states that lead into an attractor define the attractor's *basin*. For example, V, W, E, and F are part of the basin of the attractor G. That means that it does not matter from which state in that basin the process starts: the end result will always be G. This property is called *equifinality*: different initial states produce the same final state. (Note that in Newtonian mechanics there is no equifinality, because all movement is considered to be reversible.) Let us assume that the state of the system is disturbed by some outside intervention, for example pushing it out of the attractor state G into basin state F, or diverting its trajectory from E to W instead of F (see Fig. 1). As long as the deviation remains within the basin, the end result will anyway be G: the disturbance is neutralized by the system. It is as if the agent is actively intervening in order to secure the reaching of the attractor, e.g. by pulling the perturbed state back from E into G. Therefore, we can interpret G as a “goal” of the agent, i.e. a state that it desires to reach, and that it will defend against any perturbation that might push it away from this preferred state.

The trajectory that the system would follow without disturbances can be interpreted as the agent's *course of action* (Heylighen, 2012): the sequence of steps that it needs to perform in order to reach its goal starting from the present state. Each action in that sequence can be seen as intentional, directed at reaching the goal. The disturbances, which make the agent deviate from its course, on the other hand, are unintentional. They are the *challenges*, originating in the outside world, which the agent does not control, but which it may be able to tackle by appropriately changing its course of action.

This reinterpretation of a dynamical system as a goal-directed system is an application of what Dennett has called the *intentional stance* (Dennett, 1989). It assumes that the behavior of systems can be seen as if it were intentional, i.e. directed towards some future goal state. In contrast, the more traditional causal or mechanistic stance (which Dennett calls “physical”) assumes that that behavior is better seen as a sequence of causations, in which the present state produces the next state, which produces a subsequent state, and so on. As we have shown with our example—and as can be proven mathematically (Mesarović & Takahara, 1975)—, the two stances are equivalent, in the sense that the one can in principle always be translated into the other one. In practice, however, the mechanistic stance only works for simple, deterministic systems, where we can specify precisely what the next state will be given the present state. In more complex, unpredictable systems, such as organisms or social systems, on the other hand, it is easier to reason by specifying the goals, attractors or destinations that direct the overall movement of the system, because the actual trajectory will be diverted by so many unforeseeable perturbations that causal reasoning becomes essentially unreliable (Heylighen, 2012).

That is why we normally use the intentional stance to describe intelligent—typically human—agents. But there is no formal criterion to demarcate situations in which the intentional stance is appropriate from situations in which it is not. Therefore, the action ontology generalizes the intentional stance to all systems and processes. In practice, this

means that any agent defined as above in the action ontology can be characterized as having desires, intentions and sensations. These sensations are a rudimentary form of “beliefs” that the agent has about the conditions it is experiencing. This brings us to the *Beliefs-Desires-Intentions* (BDI) framework, which is a standard way of conceptualizing mind and agency (Bratman, 1999; Georgeff, Pell, Pollack, Tambe, & Wooldridge, 1999).

Psychologists have observed that people have an inborn *Theory of Mind* (ToM), based on the BDI components, which they use to predict the behavior of other people (Astington & Baird, 2005; Whiten, 1991). For example, if you know that John *desires* to see that football match, and that he *believes* that taking the bus is the best way to get from his home to the football stadium, then you can predict that he will first form the *intention* of taking that bus, and then, if no unforeseen disturbances prevent him from carrying out his intention, that he will effectively get on the bus. You can make that prediction even when you know that the road ahead is closed so that the bus will not actually reach the stadium, and therefore that John will fail to realize his desire in this way. Such ToM reasoning is an easy, efficient and natural way of predicting agents’ behaviors, even though it reduces the complex interplay of thoughts, feelings, memories and perceptions in the human mind to the simple BDI elements.

The action ontology extends this ToM/BDI conceptualization to the simplest physical agents, such as particles, molecules or bacteria. For example, you could try to predict the outcome of a ball rolling down a hill by assuming that the ball *desires* to be at the bottom of the hill, that it *believes* or *senses* that it is on a sloping surface, and that it *intends* to go down in the direction where the slope is steepest. That rudimentary model would probably give you a pretty good idea of where the ball is likely to end up, even when the hill surface is cut through by rocks, gullies and trees that constantly disturb the movement and make the actual trajectory of the ball impossible to predict.

This approach of treating physical systems as if they were intentional agents is nothing new. It is in a sense equivalent to *animism*, i.e. the belief—typical of “primitive” cultures of hunter-gatherers—that all phenomena, such as trees, animals, or mountains, are sentient beings. One advantage of an animist worldview is that it avoids alienation (Charlton, 2002; 2007), i.e. the feeling that we do not really belong to the world that surrounds us. For the alienated person of the industrial age, the environment consists of impersonal, foreign objects and mechanisms. For an animist, on the other hand, these phenomena are agents to interact with on an equal footing—as potential allies, rivals or enemies, but never as cold, inert “matter”.

Animism has been nearly universally rejected as naïve, because it anthropomorphizes simple phenomena into human-like intelligences. But the intentional stance or agent ontology does not presuppose any near-human level of intelligence: it merely attributes to all agents in-built desires, the ability to sense certain conditions, and the tendency to react to these conditions by appropriate actions. These minimal assumptions apply equally well to elementary particles and to intelligent human beings. As such, they restore a continuity and interactivity to the world that prevent us from feeling alienated from nature. Moreover, they




allow us to get rid of the mind-matter duality at the most fundamental level—and together with it a host of paradoxes and other “hard problems”.

Of course, these different agents differ radically in their level of complexity or organization. As agents become more complex and intelligent, they start to exhibit more advanced mental qualities, such as memory, feeling, reasoning or consciousness. But our underlying philosophy sees this evolution as continuous. It does not presuppose any strict boundaries between systems that exhibit these qualities (e.g. humans and higher animals) and systems that do not (e.g. insects, plants or rocks). That makes it much easier to understand the origin of complex and mysterious phenomena, such as consciousness or intelligence, by retracing their origin in much simpler phenomena.

Let us make a first step in this journey from simple to complex minds by considering systems of coordinated actions.


Organizations

The next level in the action ontology consists of networks of coupled reactions. Reactions are coupled when the output or final condition of the one is the input or initial condition of the other, like in $X \rightarrow Y$, $Y \rightarrow Z$. These couplings become more complex when input and output conditions overlap without being identical, like in:



This is typical for chemical reactions that consume and produce combinations of molecules. The metabolism of a living cell is a huge network of such coupled chemical reactions, which produce and consume a wide variety of molecules in order to provide the cell with all the energy and building blocks it needs to survive and grow.

A living organism is a typical example of a complex agent. What distinguishes such an “agent-like” network of reactions from the uncoordinated reactions that may take place e.g. in a test tube is *autopoiesis* (Maturana & Varela, 1980; Mingers, 1994; Razeto-Barry, 2012): the network produces its own components, thus maintaining an invariant organization in spite of a continuously changing state. This state is characterized by changing concentrations of molecules and a barrage of external perturbations that need to be counteracted. Autopoiesis or self-production provides the network with a stable identity in spite of the fact that it is in a situation of permanent flux. This makes it autonomous, i.e. to a significant degree independent of what happens in the environment. Still, the autopoietic network A interacts with the environment, by producing the actions Y appropriate to deal with the external challenges X . This defines the autopoietic organism as a higher-order agent:

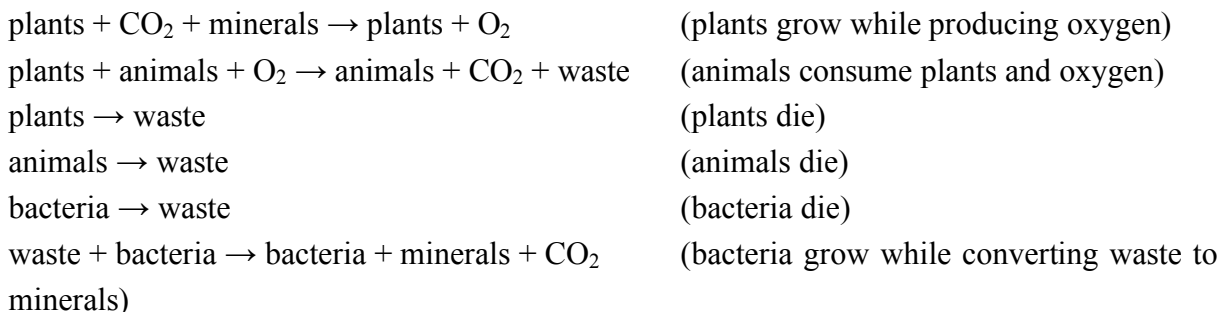


At the abstract level of this overall reaction, there is no difference between a complex agent, such as an animal or a human, and an elementary agent, such as a particle or a molecule. The difference becomes clear when we zoom in and investigate the changing state of the network of reactions inside the agent.

A very promising way to do this is the formalism of *Chemical Organization Theory* (COT) (Dittrich & Fenizio, 2007; Heylighen, Beigi, & Veloz, 2015). COT starts from reactions like the ones of the action ontology, but adds a generalized notion of autopoiesis, which it defines as the combination of closure and self-maintenance. A network characterized by closure and self-maintenance is called an *organization*. (Note that, unlike the original definition of autopoiesis, this does not include the formation of a topological boundary separating the network from its surroundings). *Closure* means that the network of reactions functions in such a way that no qualitatively new conditions are produced: in spite of all the change implied by the ongoing reactions, eventually the situation always comes back to conditions that existed before. *Self-maintenance* means that all the conditions that existed before are eventually produced again. Closure means that nothing new is created, self-maintenance that nothing old gets lost; together they imply that all the essential parts are eventually recycled. In spite of this higher-level invariance, the system is in a constant flux, as conditions are relentlessly transformed into different conditions.

Note that an organization is an attractor for the system formed by the network of reactions: through self-organization, the system sooner or later settles into a configuration that is closed and self-maintaining (Heylighen et al., 2015). Thus, autopoiesis (which is equivalent to survival and growth of the organization) is the implicit goal for a complex agent.

Perhaps the simplest example of such a self-producing organization is a cycle: $X \rightarrow Y$, $Y \rightarrow Z$, $Z \rightarrow X$. But when couplings are complex, the organization is subtler. For example, here is a highly simplified model of the ecosystem of the Earth:



All the components or conditions in this system are both consumed by some reaction and produced by some other reaction. As a result they are fully recycled: the network is closed

and self-maintaining. None of the components will ever disappear from the system, although their concentrations are constantly varying.

How does such COT model help us to understand complex agents, and in particular their cognitive or mental capabilities? COT extends the action ontology by modeling the internal processes and changing internal state of an agent. When the agent is simple, like a particle or a rock, its “belief” or “sensation” is trivial: an incoming causal signal that is directly transformed into an outgoing effect. Such agents without internal state are called *reactive* (Beer, 1995; Heylighen, 2014a): they react immediately to their sensed conditions.

With a complex agent, incoming signals (sensations) affect the internal state. The internal state thus keeps a (partial) memory of the sequence of sensations that the agent has undergone. This state is further processed by the network of internal reactions, which depends on the agent’s autopoietic organization. The resulting state may or may not result in an outgoing signal (i.e. an action affecting the outside world). This can be seen as a process of “deliberation” or “sense-making” (Stewart et al., 2014): the incoming sensation needs to be processed or interpreted, taking into account the agent’s memory of previous sensations and its implicit desire for continuing autopoiesis. It may well be that in the end it turns out that the sensation does not affect autopoiesis and that therefore no action is needed to safeguard this goal.

This is the essential difference between a complex, autopoietic agent and a simple, reactive agent. The reactive agent necessarily reacts to any condition it is sensitive to. The autopoietic agent may decide to ignore the condition, however, after having become “aware” of it and having evaluated it through its internal dynamics. One of the arguments used by Chalmers (1995) to justify why consciousness is such a “hard” problem that cannot be tackled by the traditional methods of science is that conscious experience does not seem to have a function. Indeed, we can sense, feel or experience phenomena without this affecting our actions. However, that does not mean that experience is useless: experience is the outcome of the process of sense-making, in which incoming sensations and existing memories are interpreted in terms of their implications, meaning, and valence (positive or negative) relative to our value system. Experience prepares or *primes* the mind for action that potentially needs to be performed, but that may never actually happen (Heylighen, 2014a). The zombie thought experiment views the zombie’s mind as a mechanical system in which causes (sensations) directly lead to effects (actions), while leaving no room for this complex process of sense-making that an agent needs to undergo in order to deliberate which, if any, action may be needed.

Note also that this process of deliberation, in which different alternative interpretations and possible courses of action are explored, but whose outcome is essentially unpredictable, captures our intuitive notion of “free will” (Heylighen, 2014a): many potential actions can be conceived and examined, but eventually only one (or none) is actually performed.

Socially distributed cognition

The COT model of self-sustaining networks of processes is so general that it can describe a wide variety of complex, organized systems. These include single cells, multicellular organisms, ecosystems, such as a lake or the planetary organism “Gaia”, metabolic networks, economic systems, brains, and social systems (Heylighen et al., 2015). Using the intentional stance, each of these can be described as having sensations, beliefs, desires, intentions, memories, and experiences.

Note that this list is broader than just living systems. The original definition of autopoiesis (Maturana & Varela, 1980; Varela, Maturana, & Uribe, 1974) included a requirement for the production of a physical boundary that separates the system from its surroundings. This requirement was inspired by the membrane surrounding the metabolic network that characterizes a living cell, and was intended to limit autopoiesis to living organisms. By leaving out this requirement, COT also describes self-producing networks that are distributed in space, such as markets or ecological networks. This allows it in particular to apply a clear notion of autopoiesis to social systems, a move made by several authors (e.g. Luhmann, 1986; Mingers, 1994) wanting to extend autopoiesis beyond biology.

Let us then focus specifically on the subject matter of this book and examine in how far a social system can be characterized as having some form of knowledge that is not merely the knowledge inside the brains of its human components. To clarify the matter, we first need to explain the relation between component agents and the social “super-agent” that they constitute.

COT defines a super-agent as a closed, self-maintaining network of reactions, some of which are catalyzed by simpler component agents (Heylighen et al., 2015). The network forms an “organization” when the actions of its agents are coordinated to such a degree that the whole becomes autopoietic, i.e. closed and self-maintaining, in the sense that whatever processes it undergoes, it reliably reconstitutes its own essential components. That leaves us quite some leeway in deciding which are the essential components that need to be reconstituted. Normally, these components are chosen such that they define a stable identity for the organization, meaning that subsequent states or configuration of the system can all be identified as aspects of the same organization.

In Luhmann’s theory of autopoietic social systems (Luhmann, 1986, 1995), these components are distinctions or symbols that are processed via communications from agent to agent within the social system, but such that the process leaves the essential organization invariant. This notoriously difficult theory can actually be formalized—at least in part—rather easily by means of COT (Dittrich & Winter, 2005).

In our own approach to social systems, we model social processes through the *propagation of challenges* (Heylighen, 2014b). This is a generalization of Hutchins’s analysis of socially distributed cognition taking place through the propagation of “state” (Hutchins, 1995, 2000): the state of some agent determines that agent’s action or communication, which

in turn affects the state of the next agent receiving that communication or undergoing that action. Since a state is a specific selection out of a variety of potential states, it carries information. Therefore, the propagation of state from agent to agent is equivalent to the transmission and processing of information. This is an acceptable model of distributed cognition if cognition is conceived as merely complex information processing. But if we want to analyze cognition as the functioning of a mind or agency, then we cannot ignore that agent's desire, or more broadly its system of values and preferences. What counts for an agent is not so much the objective state of some phenomenon, but the degree to which that state affects the agent's values: in how far does it either help or hinder the agent in realizing its desires? This shifts our view of information from the traditional syntactic perspective of information theory (information as selection among possibilities (Shannon & Weaver, 1963)) to a pragmatic perspective (information as trigger for goal-directed action (Gernert, 2006)).

We make this change of perspective more concrete by replacing the terms "information" or "state" by "challenge". A *challenge* is defined as a situation (i.e. a conjunction of conditions or states perceived by some agent) that stimulates the agent to act, because acting on that challenge would bring benefit to the agent relative to not acting (Heylighen, 2012). Challenges can be positive (acting brings the agent closer to its goal) or negative (not acting pushes the agent farther away from its goal). Positive challenges can be seen as opportunities for advancing towards the goal or as resources to be exploited; negative challenges as problems to be resolved or dangers to be evaded. For example, a tasty treat is a positive challenge that will elicit the action "eat". A poisonous snake is a negative challenge that will elicit the action "run away".

By acting on a challenge, an agent will change the situation. If the challenge is fully "relaxed" (opportunity exploited or problem solved) (Heylighen, 2014b), then the new situation will no longer be a challenge. However, for complex challenges, such as building a house, a single agent can in general not fully resolve it. For example, an architect may make a plan for the house, but cannot build the house without help from others. In this case, the new situation (available plan) constitutes a challenge for one or more agents (e.g. contractor, builders, carpenter, plumber...) to perform the implied next actions. After each action, the situation moves closer to a full resolution of the initial challenge. Yet, so long as that end has not been reached the resulting situation defines a new challenge. Thus, challenges propagate from agent to agent until full relaxation.

This propagation typically follows the connections within an organization or social network, as people pass on challenges to subordinates, collaborators or friends. In formally structured organizations, such as an administration, a building firm or a factory, such propagation follows predefined paths, called *workflows* (Van der Aalst & Van Hee, 2004), in which a complex task is decomposed into a number of more specific tasks to be executed in a particular sequence by agents performing specialized roles. Each individual contributes his or her specific skills or expertise to tackling some part of the challenge. Thus, knowledge about how to solve the problem is divided across a variety of agents, some of which (e.g. computer

programs, robots, measuring apparatuses, manuals with formal guidelines, ...) may not be human. But knowledge and processing is distributed across more than individual agents: it is distributed across the network of actions that connects them.

This can be understood by going back to the elements of the action ontology: actions represented as production rules of the form $X \rightarrow Y$. Each agent can be characterized by a collection of production rules representing the actions that agent is able to perform. This includes both actions that change the external situation and internal, “mental” actions that constitute the process of sense-making, in which the agent interprets the incoming information and eventually formulates a course of (external) action. These cognitive actions can be seen as *inferences* in which some condition Y is inferred from some previously established condition X. Here is an example of a three-step inference, starting from the perception of a snake, and ending with the intention to flee:

snake \rightarrow poisonous
poison \rightarrow mortal danger
mortal danger \rightarrow flee

Seen from the outside, the agent behaves as if it follows the single rule: snake \rightarrow flee. The intermediate inferences are not directly observable. They implement a (very simple) process of deliberation, in which the perception is assessed against pre-existing beliefs (that snakes can be poisonous, and that a bite from a poisonous snake can kill) and desires (for survival) in order to decide about a course of action (that it is best to flee).

Each agent in a workflow will use a variety of such internal inference rules to process the incoming challenge and form a plan of action. It will then execute the plan while monitoring the provisional results, and if necessary, use this feedback to correct the course of action until it achieves the intended goal. Assuming that these inference rules are correct and justified, they can be seen as the agent’s *knowledge* about how to tackle this kind of challenges. The question now is whether a social system possesses knowledge that is not located in the brain of its human components.

The case is most obvious for rules that are implemented in material supports, such as calculators or computer programs, but also documents listing rules too numerous or complicated for anyone to remember. For example, if tackling the challenge at some stage requires the calculation of the logarithm of a number, no human person will calculate that logarithm inside the head, but rather enter the number into a calculator and register the result. Before computing technology, that same person would have searched for the number in a big book with logarithm tables, and similarly noted down the corresponding result. In neither case would we have found any person in the organization who knows the logarithms for all relevant numbers, i.e. whose brains would have contained the production rules of the form $\log(x) \rightarrow y$, for any substantial series of numbers x. However, the organization as a whole does know how to calculate a logarithm. The same applies for various non-mathematical

operations that an organization regularly performs. These tend to be written down in the form of manuals, guidelines, regulations, or procedures. These documents may have to be consulted in order to ascertain that a specific condition X is a special case of condition Y, which requires the performance of action Z, while checking for the presence of condition W, and so on.

But the case for distributed knowledge can be made even in a purely social system, where all information processing is done through individual thinking and communication between individuals. The situation is perhaps most intuitive for procedural knowledge. A complex item, such as a car or a computer, is never manufactured by a single person. No person knows how to perform all the operations that are necessary to assemble all the components of a car. But together, all the employees of a car factory can build a car, because each one will apply his or her specific skill by adding, adjusting or assembling these particular components. Nevertheless, it is not sufficient to gather a group of individuals having each of those skills in one big room for them to start building a car. The workflow or process itself is crucial, because these specialized skills can only be applied at the precise moment when all the preparatory actions have been performed. Thus, person B needs to know that some person A has performed the action $X \rightarrow Y$, before B can perform the subsequent action $Y \rightarrow Z$. The procedural knowledge of how to build a car is more than the aggregate of the procedures that the different employees have in their brain: the workflow connecting these individuals and their procedural knowledge to each other is itself part of the overall procedure.

This argument can be extended to declarative or semantic knowledge. Suppose that John receives a phone call telling him about condition U (say a specific client request). He knows that this is a special case of condition V, but otherwise does not know anything about V. Therefore, he passes on the challenge to his colleague Ann, whom he assumes to be more knowledgeable about this domain. Ann knows that a crucial part of V is W, and that her colleague Barbara is an expert in W-related matters. Barbara immediately sees that W entails X, a type of problem dealt with by the people of the X-matters department, including Tom. Tom recommends tackling X by action Y, which is finally executed by Jane, thus satisfying the client. In this way, the organization to which John, Ann, Barbara, Tom and Jane belong has performed a process of inference of which the different steps are distributed across the different agents:

John: $U \rightarrow V$

Ann: $V \rightarrow W$

Barbara: $W \rightarrow X$

Tom: $X \rightarrow Y$

The process can be summarized by saying that the organization knows that U entails X and is to be dealt with by action Y. But no single agent in the organization actually knows that $U \rightarrow X$. It is not sufficient that each of the rules leading to that conclusion is known by some

individual in the organization for that conclusion to be effectively drawn. The agents must moreover be organized in such a way that the next step in the process is passed on to an agent who knows how to perform this step. Thus, the network of relations within the organization is an essential part of that organization's knowledge of how to deal with challenges. Therefore, the knowledge must be conceived as socially distributed.

While we have here focused on objective knowledge (U entails X), the same reasoning can be made about more subjective mental phenomena such as perception, meaning, awareness or desire. In a typical process of challenge propagation in an organization, decisions are made based on the feelings, values and desires of the different agents along the propagation chain. Together, these determine the overall course of action of the organization. This course of action moves towards a certain "attractor", which defines the collective desire of the organization. While moving in this direction, the organization continues to collect information about its situation through a variety of channels involving different agents, while trying to make sense of that information, and deliberating whether this awareness of the situation requires some change in its course of action.

Experiencing non-duality

We have argued that both simple physical agents and (self-)organized networks of processes can be conceptualized as mind-like agencies. Thus, mind does not just reside inside the brain; it is distributed across the whole of nature and society. This implies a radical negation of the mind-matter duality: not only is it impossible to have a mind independent of matter, it is impossible to find matter that does not exhibit some kind of mind-like properties.

Although the above argument may appear logically coherent, it is unlikely to be convincing on a more intuitive level. After all, we all feel that our mind is sitting somewhere inside our skull, looking out at the external world, experiencing its sensations, and pondering what to do next—don't we? Apart from other human beings such as us (and perhaps some of the smarter animals) that outside world lacks intelligence, feeling, desire, or awareness—doesn't it? That world is merely a collection of inert, material objects and mechanisms, ready to be manipulated through the actions conceived by our independent mind—isn't it?

While these intuitions may be common in our materialistic and reductionistic, Western society, there are in no way universal. We already noted that prehistoric thought was fundamentally animistic (Charlton, 2007). Moreover, Eastern civilizations have produced a number of holistic philosophies, such as Taoism, Buddhism and to some extent Sufism, that emphasize non-duality, advising people to give up the illusion of the self as a separate agent, and to seek reconnection with the world. This world is conceived as an immense process or flow, a "Becoming" or "Tao". Individuals should not try to control this flow, the way Western science and technology try to control nature, but go along with it, by becoming aware that the

self is fluid and that it does not have any clear boundaries with the surrounding flow of existence.

Although Taoism's picture of the inseparability between mind and world appears alien when seen through the lens of analytic philosophy, it provides a cue for reconceptualizing how science can solve the complex challenges of our time. Take for instance the grand challenge of urbanization and the need for developing sustainable and resilient communities. Formulating the connection between our actions, our mental models and the surrounding world through non-dual philosophies can help us to develop engaging narratives and models for change. *Mindfulness Engineering* (Beigi, 2014), for example, is an approach that integrates holistic approaches to the mind-body connection with the engineering of sustainable and resilient cities. In this view, cities are not just aggregates of buildings, roads and other material infrastructures, but organized networks of people and objects that together constitute a “super-agent”. Only through understanding the nature of the human mind and its connection with other minds and the physical environment can we design a truly “smart” city (Chourabi et al., 2012), which is able to deal with complex challenges such as pollution, traffic jams or earthquakes. Part of the solution is to engineer the environment in such a way that it “nudges” people to behave in more desirable ways (Heylighen, Kostov, & Kiemen, 2013; Thaler & Sunstein, 2008), thus using the world to affect the mind which in turn affects the world. This illustrates how pragmatic engineering and technical solutions inspired by distributed cognition and action ontology can revitalize old and unsustainable ways of doing things.

The complex, dynamic and interconnected nature of the social, technological and ecological systems that surround us is now well recognized by scientists (e.g. Ball, 2012; Helbing, 2012; Walker, Holling, Carpenter, & Kinzig, 2004). However, it is still too often ignored by the actual decision-makers—the politicians, managers or employees—who need to develop a sustainable course of action that takes into account the interdependency of our world. Scientific theories, such as systems dynamics, complex adaptive systems and distributed cognition, especially when supported by more general philosophical frameworks, such as process metaphysics or action ontology, should help people to better understand and deal with these complexities.

Still, our decision-making is ideally supported not only by theoretical understanding, but by the concrete experience of connectedness, non-duality and flow. This is what the practices of Buddhism and Taoism try to achieve, through techniques such as meditation, yoga and tai chi. For example, intense meditation can result in a so-called “oceanic feeling” in which the subject no longer feels like a separate individual, but as merging with the larger whole. Different spiritual practices aim at achieving a more enduring state of “enlightenment” (Harris, 2014), in which people no longer sense the need to control their situation by cautiously planning all their actions, but instead are able to act spontaneously, without worry or rumination, while being in “flow” (Nakamura & Csikszentmihalyi, 2002) with the process, and “mindful” (Bishop et al., 2004) of their surroundings. Such an altered state of

consciousness could be seen as a form of recovered animism, in which the individual again feels part of an encompassing network of interactions (Charlton, 2007).

Western observers who attempt these practices typically emphasize the great effort, discipline, and time needed to achieve any form of non-dual consciousness (Harris, 2014). For most people, it seems very difficult to get rid of the impression that there is some individual self sitting inside the Cartesian theater, talking to itself, and looking out at the world. However, from our perspective, this difficulty appears like an artefact of our upbringing. Small children cannot yet distinguish between self and world (Rochat, 2001), and need to learn that some movement they perceive is the result of their own action rather than an outside event. The concept of “object” only arises after the subject has learned to distinguish between purely internal actions (e.g. moving the head in such a way that the image on the retina changes) and actions that affect the outside world (e.g. manipulating a toy). Older children’s more elaborate concept of self is to an important degree a product of society, which teaches them that they are individuals with a particular identity, role and duties towards the rest of society. These duties are interiorized in the form of what Freud has called the “super-ego”. For many people, this appears like a little voice inside their head that is constantly reminding them of the rules they should be heeding, and deliberating verbally which is the best action to take. Most people find it very difficult to shut off that relentless inner monologue (Harris, 2014), and therefore to experience the world as it originally appears to the senses, i.e. as a continuous, non-verbal flow of interactions.

Yet, for some people such non-dual, mindful awareness seems to be their default state, while their discursive self is merely a cognitive tool that they switch on whenever the situation requires explicit reflection rather than the intuitive going with the flow. We have recently started to investigate this condition under the label of “meta-awareness”. It illustrates that non-dualism is not just a theoretical position, but a concrete attitude and experience that can help us to cope with a complex and dynamic world.

Conclusion

We have approached the problem of the extended mind, and in particular of socially extended knowledge, from a radically non-dualist perspective. Mind-matter dualism is an artefact of the Newtonian worldview, which reduces all phenomena to the mechanical motion of material bodies, governed by deterministic laws. In this picture, there is no room for free will, agency, desire, sensation or experience. Therefore, both philosophers and laypeople are inclined to situate these phenomena in the distinct, non-material realm of mind, even when they believe that this mind still somehow supervenes on matter. This artificial and inconsistent separation between mind and matter creates a host of apparently unanswerable questions, paradoxes and other “hard problems”. For us, the only way out is to get rid of dualism at the most fundamental level.

We have proposed to do that by introducing an ontology of action, which can be seen as a concrete, scientifically underpinned implementation of a process philosophy. The elements of this ontology are actions or reactions. These have the form $X \rightarrow Y$, representing an elementary process that leads from some condition X to a new condition Y . Agents are defined as catalysts of such reactions, i.e. conditions necessary for the reaction to take place, but that are not themselves affected by the reaction. The different reactions triggered by an agent A constitute the actions that A is capable of executing. We then applied the intentional stance by interpreting an agent's actions as goal-directed. This makes sense because these actions are characterized by equifinality: they lead from a variety of initial conditions (the basin) to the same end condition (the attractor), and thus implicitly resist disturbances that make them deviate from this trajectory.

That allowed us to characterize an agent as having sensations or *beliefs* (the conditions to which the agent reacts), *desires* (the attractors the agent tries to reach), and *intentions* (its expected course of action leading towards an attractor). This BDI conceptualization fits in with the “theory of mind” that people intuitively use to predict and explain the behavior of others. Its extension to the simplest kind of agents explains why in pre-scientific cultures mind and agency are so easily ascribed to non-human phenomena—a way of thinking known as animism. Animism has been abandoned as a description of physical phenomena because these can in general be described more precisely through a causal model. However, an intentional description can in principle be made equivalent with a description in terms of cause and effect. Moreover, it has the advantage that it allows predictions and explanations even in cases where the situation is too complex for mechanistic modeling.

We have proposed to represent such more complex phenomena as networks of reactions and agents. Such networks tend to self-organize to a configuration where they become self-maintaining or autopoietic. That means that they develop an invariant identity within a flux of endless change by continuously rebuilding their essential components. This property can be elegantly expressed in the action ontology with the help of the formalism of Chemical Organization Theory (Dittrich & Fenizio, 2007; Heylighen et al., 2015). It turns the network into a higher-order autonomous system: a *super-agent*. In contrast to an elementary, “reactive” agent, such a complex agent has an internal state that is affected by present and previous input or “sensations”, thus keeping some sort of memory of these interactions. This state in turn affects the agent's output of “actions”. The intermediate process, where a host of sensations, memories, and internal, goal-directed dynamics interact to produce a continuously changing state, can be seen as the agent's process of “sense-making” and “deliberation”, in which it interprets the situation and potentially decides about a course of action.

This general, abstract model of a super-agent can be applied to social systems. Here the components agents are people and their material supports, such as books or computers. The process of deliberation can here be seen as a form of *distributed cognition*: the different human and technological agents interpret the situation, make inferences, solve problems, and plan actions by propagating information, or more precisely *challenges*, from the one to the

other, along the links in the social or organizational network. Each agent in such a workflow will typically contribute its own specialized knowledge to tackling the challenge. However, in general the outcome is emergent: no individual agent has the knowledge to deduce the final solution from the initial problem. That knowledge is distributed, not only across human individuals and their external memories, but across the links in the organizational network: the same agents linked in a different way may not be able to collectively solve the problem.

We have concluded our review of the action ontology perspective on mind by examining why such a philosophy, in spite of its simplicity and coherence, is unlikely to be easily accepted. The Cartesian split between mind and matter is a basic tenet of our Western culture, with its attitude of materialism and mechanicism, on the one hand, and of individual freedom and autonomy, on the other hand. The more holistic and collectivistic Eastern cultures are more prone to see the world as a single, indivisible process, in which the separation between self and world is an illusion. They even propose concrete techniques, such as meditation, to help people free themselves from this illusion. But the fact that these disciplines are typically experienced as very demanding shows that this “illusion” has a pretty strong psychological basis.

Yet, there is evidence that the self–world distinction is not as pervasive as generally assumed, and that people can spontaneously experience themselves as part of an encompassing flow rather than as an independent mind in its Cartesian theater. Such a non-dual awareness is worth promoting, because it not only seems to prevent alienation and other sources of psychological suffering (Charlton, 2007; Harris, 2014), but helps people to better understand how they fit in with the complex processes that surround them.

The action ontology may support this general enterprise of raising awareness of the inseparability of mind and world by integrating the broad, but vague, outlook of process metaphysics and Eastern philosophy with the clarity and precision of more analytic, formal models of actions and networks.

Acknowledgements

This research was supported by the Yuri Milner Foundation as part of the Global Brain Institute.

References

- Anderson, M. L. (2003). Embodied cognition: A field guide. *Artificial intelligence*, 149(1), 91–130.
- Astington, J. W., & Baird, J. A. (Eds.). (2005). *Why language matters for theory of mind*. New York: Oxford University Press.

- Ball, P. (2012). *Why Society is a Complex Matter: Meeting Twenty-first Century Challenges with a New Kind of Science* (2012th ed.). New York: Springer.
- Beer, R. D. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72(1-2), 173–215.
- Beer, R. D. (2000). Dynamical approaches to cognitive science. *Trends in cognitive sciences*, 4(3), 91–99.
- Beigi, S. (2014). *Mindfulness Engineering: A Unifying Theory of Resilience for Volatile, Uncertain, Complex and Ambiguous (VUCA) World* (PhD thesis). University of Bristol, Bristol, UK.
- Bishop, S. R., Lau, M., Shapiro, S., Carlson, L., Anderson, N. D., Carmody, J., Segal, Z. V., et al. (2004). Mindfulness: A Proposed Operational Definition. *Clinical Psychology: Science and Practice*, 11(3), 230–241. doi:10.1093/clipsy.bph077
- Bratman, M. E. (1999). *Intention, Plans, and Practical Reason* (New edition ed.). Stanford, Calif: The Center for the Study of Language and Information Publications.
- Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3), 200–219.
- Charlton, B. (2002). What is the Meaning of Life? Animism, Generalised Anthropomorphism and Social Intelligence. Retrieved from <http://www.hedweb.com/bgcharlton/meaning-of-life.html>
- Charlton, B. G. (2007). Alienation, recovered animism and altered states of consciousness. *Medical Hypotheses*, 68(4), 727–731. doi:10.1016/j.mehy.2006.11.004
- Chourabi, H., Nam, T., Walker, S., Gil-Garcia, J. R., Mellouli, S., Nahon, K., Pardo, T. A., et al. (2012). Understanding Smart Cities: An Integrative Framework. *2012 45th Hawaii International Conference on System Science (HICSS)* (pp. 2289–2297). Presented at

- the 2012 45th Hawaii International Conference on System Science (HICSS).
doi:10.1109/HICSS.2012.615
- Clark, A. (1998). Embodied, situated, and distributed cognition. *A companion to cognitive science*, 506–517.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7–19. Retrieved from <http://cogprints.org/320/1/extended.html>
- Dennett, D. C. (1989). *The intentional stance*. The MIT Press.
- Dennett, D. C., & Kinsbourne, M. (1997). Time and the observer: The where and when of consciousness in the brain. *The Nature of Consciousness: Philosophical Debates*, 141–174.
- Dittrich, P., & Fenizio, P. S. di. (2007). Chemical Organisation Theory. *Bulletin of Mathematical Biology*, 69(4), 1199–1231. doi:10.1007/s11538-006-9130-8
- Dittrich, P., & Winter, L. (2005). Reaction networks as a formal mechanism to explain social phenomena. *Proc. Fourth Int. Workshop on Agent-Based Approaches in Economics and Social Complex Systems (AESCS 2005)* (pp. 9–13). Retrieved from <http://users.minet.uni-jena.de/~dittrich/p/DW2005.pdf>
- Georgeff, M., Pell, B., Pollack, M., Tambe, M., & Wooldridge, M. (1999). The belief-desire-intention model of agency. *Intelligent Agents V: Agents Theories, Architectures, and Languages* (pp. 1–10). Springer. Retrieved from http://link.springer.com/chapter/10.1007/3-540-49057-4_1
- Gernert, D. (2006). Pragmatic Information: Historical Exposition and General Overview. *Mind and Matter*, 4(2), 141–167.
- Harris, S. (2014). *Waking Up: A Guide to Spirituality Without Religion*. Simon and Schuster.

- Helbing, D. (2012). Managing Complexity. In D. Helbing (Ed.), *Social Self-Organization, Understanding Complex Systems* (pp. 285–299). Springer Berlin Heidelberg. Retrieved from http://link.springer.com/chapter/10.1007/978-3-642-24004-1_15
- Heylighen, F. (1999). The growth of structural and functional complexity during evolution. In F. Heylighen, J. Bollen, & A. Riegler (Eds.), *The evolution of complexity* (pp. 17–44). Springer. Retrieved from <http://pcp.vub.ac.be/Papers/ComplexityGrowth.pdf>
- Heylighen, F. (2011). Self-organization of complex, intelligent systems: an action ontology for transdisciplinary integration. *Integral Review*. Retrieved from <http://pespmc1.vub.ac.be/Papers/ECCO-paradigm.pdf>
- Heylighen, F. (2012). *A Tale of Challenge, Adventure and Mystery: towards an agent-based unification of narrative and scientific models of behavior* (ECCO Working Papers No. 2012-06). ECCO. Brussels, Belgium. Retrieved from <http://pcp.vub.ac.be/papers/TaleofAdventure.pdf>
- Heylighen, F. (2014a). *Cognitive Systems: a cybernetic approach on the new science of the mind* (Lecture Notes.). ECCO, VUB. Retrieved from <http://pespmc1.vub.ac.be/Papers/CognitiveSystems.pdf>
- Heylighen, F. (2014b). Challenge Propagation: Towards a theory of distributed intelligence and the global brain. *Spanda Journal*, *V*(2), 51–63. Retrieved from <http://cleamc11.vub.ac.be/papers/ChallengePropagation-Spanda.pdf>
- Heylighen, F., Beigi, S., & Veloz, T. (2015). *Chemical Organization Theory as a universal modeling framework for interaction, self-organization, and autopoiesis* (ECCO Working Papers). Retrieved from <http://pespmc1.vub.ac.be/Papers/COT-applicationsurvey.pdf>
- Heylighen, F., Busseniers, E., Veitas, V., Vidal, C., & Weinbaum, D. R. (2012). *Foundations for a Mathematical Model of the Global Brain: architecture, components, and*

- specifications* (GBI Working Papers No. 2012-05). GBI Working Papers. Retrieved from <http://pespmc1.vub.ac.be/papers/TowardsGB-model.pdf>
- Heylighen, F., Cilliers, P., & Gershenson, C. (2007). Complexity and Philosophy. In J. Bogg & R. Geyer (Eds.), *Complexity, science and society* (pp. 117–134). Oxford: Radcliffe Publishing,. Retrieved from <http://arxiv.org/abs/cs/0604072>
- Heylighen, F., Kostov, I., & Kiemen, M. (2013). Mobilization Systems: technologies for motivating and coordinating human action. In Peters M. A., Besley T. and Araya D. (Ed.), *The New Development Paradigm: Education, Knowledge Economy and Digital Futures*. Routledge. Retrieved from <http://pcp.vub.ac.be/Papers/MobilizationSystems.pdf>
- Hutchins, E. (1995). *Cognition in the Wild* (Vol. 262082314). Cambridge, MA: MIT press. Retrieved from http://books.google.be/books?id=AfupQgAACAAJ&dq=Cognition+in+the+Wild&hl=en&sa=X&ei=YjLWT6PGLaOn0QWwyqiBBA&redir_esc=y
- Hutchins, E. (2000). Distributed cognition. In N. J. Smelser & P. B. Baltes (Eds.), *International Encyclopedia of the Social and Behavioral Sciences*. Elsevier Science.
- Luhmann, N. (1986). The autopoiesis of social systems. *Sociocybernetic paradoxes*, 172–192. Retrieved from <http://books.google.be/books?hl=en&lr=&id=2Yfys20NPAoC&oi=fnd&pg=PA172&dq=Luhmann&ots=E5hsDIDSAG&sig=rGBVijKWKxrLU2NTkVWizu6rZko>
- Luhmann, N. (1995). *Social systems*. Stanford University Press. Retrieved from http://books.google.be/books?hl=en&lr=&id=zVZQW4gxXk4C&oi=fnd&pg=PR9&dq=Luhmann&ots=7EEKk78LVS&sig=GNMqh5YJn54eoK_BQzuRzrgFdOI
- Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and Cognition: The Realization of the Living*. D Reidel Pub Co.

- McGinn, C. (1989). Can We Solve the Mind–Body Problem? *Mind*, *XCVIII*(391), 349–366.
doi:10.1093/mind/XCVIII.391.349
- Mesarović, M. D., & Takahara, Y. (1975). *General systems theory: mathematical foundations* (Vol. 113). Elsevier Science.
- Mingers, J. (1994). *Self-Producing Systems: Implications and Applications of Autopoiesis*. Springer Science & Business Media.
- Nakamura, J., & Csikszentmihalyi, M. (2002). The concept of flow. In C. R. Snyder (Ed.), *Handbook of positive psychology* (pp. 89–105). New York, NY: Oxford University Press.
- Razeto-Barry, P. (2012). Autopoiesis 40 years Later. A Review and a Reformulation. *Origins of Life and Evolution of Biospheres*, *42*(6), 543–567. doi:10.1007/s11084-012-9297-y
- Rescher, N. (1996). *Process metaphysics: An introduction to process philosophy*. State Univ of New York Pr.
- Rochat, P. (2001). Origins of self-concept. *Blackwell handbook of infant development*, *2*, 191–212.
- Seager, W. (2006). Panpsychism. *Encyclopedia of Cognitive Science*. John Wiley & Sons, Ltd. Retrieved from <http://onlinelibrary.wiley.com.ezproxy.vub.ac.be:2048/doi/10.1002/0470018860.s00119/abstract>
- Shannon, C. E., & Weaver, W. (1963). *The Mathematical Theory of Communication*. University of Illinois Press.
- Stewart, J., Gapenne, O., & Paolo, E. A. D. (Eds.). (2014). *Enaction: Toward a New Paradigm for Cognitive Science* (Reprint ed.). A Bradford Book.

- Susi, T., & Ziemke, T. (2001). Social cognition, artefacts, and stigmergy: A comparative analysis of theoretical frameworks for the understanding of artefact-mediated collaborative activity. *Cognitive Systems Research*, 2(4), 273–290.
- Teilhard de Chardin, P. (1959). *The phenomenon of man*. Collins London.
- Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale Univ Pr.
- Toulmin, S. (1993). From Clocks to Chaos: Humanizing the Mechanistic World-View. In H. Haken, A. Karlqvist, & U. Svedin (Eds.), *The Machine as Metaphor and Tool* (pp. 139–153). Berlin, Heidelberg: Springer Berlin Heidelberg. Retrieved from http://www.springerlink.com/index/10.1007/978-3-642-77711-0_9
- Turchin, V. F. (1993). The Cybernetic Ontology of Action. *Kybernetes*, 22(2), 10–30. doi:10.1108/eb005960
- Van der Aalst, W., & Van Hee, K. M. (2004). *Workflow management: models, methods, and systems*. The MIT press.
- Varela, F. J., Maturana, H. R., & Uribe, R. (1974). Autopoiesis: the organization of living systems, its characterization and a model. *Biosystems*, 5(4), 187–196. Retrieved from <http://www.sciencedirect.com/science/article/pii/0303264774900318>
- Walker, B., Holling, C. S., Carpenter, S. R., & Kinzig, A. (2004). Resilience, adaptability and transformability in social–ecological systems. *Ecology and society*, 9(2), 5. Retrieved from ftp://131.252.97.79/Transfer/WetlandsES/Articles/walker_04_socio-ecology_resilience.pdf
- Whitehead, A. N. (1978). *Process and Reality: An Essay in Cosmology* Ed. David Ray Griffin and Donald W. Sherburne. Free Press, New York.[PR].
- Whiten, A. (Ed.). (1991). *Natural theories of mind: Evolution, development and simulation of everyday mindreading*. Cambridge, MA, US: Basil Blackwell.

