

Relational agency

Heylighen, Francis

Published in:
Evolution 'On Purpose'

DOI:
[10.7551/mitpress/14642.003.0008](https://doi.org/10.7551/mitpress/14642.003.0008)

Publication date:
2023

Document Version:
Accepted author manuscript

[Link to publication](#)

Citation for published version (APA):
Heylighen, F. (2023). Relational agency: a new ontology for co-evolving systems. In P. Corning (Ed.), *Evolution 'On Purpose': Teleonomy in Living Systems* (pp. 79-104). (Vienna Series in Theoretical Biology). MIT Press.
<https://doi.org/10.7551/mitpress/14642.003.0008>

Copyright

No part of this publication may be reproduced or transmitted in any form, without the prior written permission of the author(s) or other rights holders to whom publication rights have been transferred, unless permitted by a license attached to the publication (a Creative Commons license or other), or unless exceptions to copyright law apply.

Take down policy

If you believe that this document infringes your copyright or other rights, please contact openaccess@vub.be, with details of the nature of the infringement. We will investigate the claim and if justified, we will take the appropriate steps.

Relational agency: a new ontology for co-evolving systems

Francis Heylighen
Center Leo Apostel
Vrije Universiteit Brussel

Abstract: A wide variety of approaches and mechanisms have been proposed to “extend” the neo-Darwinist theory of evolution, including self-organization, symbiogenesis, teleonomy, systems biology and niche construction. These extensions share a focus on agents, networks and processes rather than on independent, static units, such as genes. To develop a new evolutionary synthesis, we therefore need to replace the traditional object-based ontology by one that is here called “relational agency”. The paper sketches the history of both object-based and relational worldviews, going back to their roots in animism and Greek philosophy. It then introduces the basic concepts of the relational agency model: condition-action rules, challenges, agents, reaction networks and chemical organizations. These are illustrated with examples of self-contained ecosystems, genes and cells. The fundamental evolutionary mechanism is that agencies and reactions mutually adapt so as to form a self-maintaining organization, in which everything consumed by one process is produced again by one or more other processes. Such autonomous organization defines a higher-level agency, which will similarly adapt, and thus become embedded in a network of relationships with other agencies.

Introduction

The dominant paradigm in biology since about the mid-twentieth century has been neo-Darwinism—also known as the “modern synthesis” of Darwinian evolution with genetics (Mayr and Provine 1980). It assumes that all features of living systems can be explained by evolution at the level of the genes. More precisely, the theory reduces evolution to changes in the frequency of certain genes in the population under the influence of natural selection. Selection is here understood as the differential reproduction of organisms depending on how well their genetic program is adapted to the environment in which they live. In other words, the environment imposes a “selective pressure” or “force of natural selection” that increases the frequency of alleles that confer higher fitness, while reducing the frequency of the others. From the neo-Darwinian perspective, the only additional mechanism needed to explain

evolution is random variation that creates new, potentially fitter genes, from which the best ones can then be selected, in an on-going process that generates organisms that are ever better adapted to their environment.

Neo-Darwinism has been very successful in proposing a simple, clear and explicit mechanism that can explain a wide range of biological phenomena, many of which seemed difficult to explain otherwise—such as the precarious balance between selfish and cooperative behaviors (Dawkins 2006). However, over the last decades its assumptions have drawn increasing criticism for being too reductionist, and in particular for neglecting a wide range of phenomena characterized by complex interactions. These interactions take place at, and in between, different levels of complexity, including molecules, genes, organelles, cells, tissues, organs, organisms, groups, ecosystems and physical environments. Of these, neo-Darwinism only really focuses on genes and environments, while at best acknowledging the importance of organisms as the phenotypes, or gene-carrying “vehicles”, that mediate between the two (Dawkins 2006).

A wide variety of approaches have tried to complement or “extend” the evolutionary synthesis with some of these more complex processes (Fábregas-Tejeda and Vergara-Silva 2018a; Laland et al. 2015; Pigliucci and Müller 2010). There are too many of those to review in the space of this introductory section, so I will here just list the most influential ones.

Self-organization is a mechanism that explains how complex organization can emerge from local interactions without need for selection by the environment (Ashby 1962; Camazine et al. 2003; Heylighen 2001; Kauffman 1993). *Systems biology* is an approach that focuses on the complex network of reactions, pathways and cycles necessary to maintain a living system’s metabolism (Palsson 2015). *Co-evolution* is the idea that the environment of an organism mostly consists of other organisms, each of which is adapting to the changes occurring in the others. Thus, they are pushing each other to evolve further and further. *Synergy* is the idea that the resulting relations between organisms are not just competitive, but that they can be mutually beneficial, creating selection for synergetic relationships (Corning 2003). *Symbiogenesis* is the formation of more complex organisms by the merging of independently evolved organisms that have developed some form of synergy (Agafonov, Negrobov, and Igamberdiev 2021). *Niche construction* notes that organisms are not just adapting to a pre-existing environment, but that they themselves change the environment in which they live, thus affecting their own selection (K. N. Laland, Matthews, and Feldman 2016). *Teleonomy* refers to the observation that development and behavior of organisms is not just passively reacting to the environment, but actively directed towards remote goals, by compensating for perturbations that otherwise would make the system deviate from these goals (Mayr 1974; P. A. Corning 2022). *Evo-devo* is an approach that focuses on how the ontogenetic development of embryos into complex organisms both constrains and enables the organic forms that can evolve on the genetic level (Fábregas-Tejeda and Vergara-Silva 2018b; Müller 2007). *Multi-level selection* notes that selection takes place not only at the level of genes, but also at the

levels of cells, colonies, multicellular organisms and groups of organisms (Okasha 2005; Wilson 1997). The theory of *major evolutionary transitions* tries to understand how cooperative interactions between systems at one level, such as cells, can give rise to the emergence of a collective system at a higher level, such as a multicellular organism (Maynard Smith and Szathmary 1997; West et al. 2015).

While seemingly disparate, there are two assumptions these different approaches have in common. First, they are *relational* (Heylighen 1990; Marmodoro and Yates 2016). Instead of reducing biological systems to independent units, such as genes or individuals, they investigate how these units interact and thus form part of a network of interdependencies that may give rise to a higher-order system. Second, they emphasize autonomous action or *agency*. Instead of seeing organisms as passively undergoing the forces of natural selection, they investigate how systems and processes actively seek or produce fit, synergetic arrangements, while counteracting environmental influences that push them away from these preferred states. Putting the two assumptions together leads me to formulate a new, synthetic perspective that I will call *relational agency*. Simply put, this is an approach that sees the world as a network of interacting agencies rather than a collection of independent objects subjected to external forces. In the remainder of this paper, I will try to formulate some foundational concepts and principles for this worldview.

Yet, to really understand the revolutionary implications of this perspective, we first need to distinguish it from the more traditional object-based worldview, whose ontology Walsh has summarized as “objectcy” (Walsh 2018). This ontology is implicit not only in neo-Darwinism, but in most modern scientific theories. As I hope to show, these implicit assumptions about the nature of reality make it intrinsically difficult to go beyond the traditional perspective and develop a true understanding of complex interactions. Understanding them better is a necessary first step in developing an alternative ontology that transcends the limitations of the object-based view.

The ontology of objects

The modern scientific worldview originated and developed within the European culture of the last 2500 years, whose roots go back to Greek philosophy. Parmenides, supported by his student Zeno, argued that change (“Becoming”) is an illusion, and that only eternal “Being” is real. Inspired by the abstract structures of geometry, Plato postulated that the ever-varying phenomena we perceive are merely imperfect reflections of absolute, unchanging forms. While his student Aristotle was more interested in concrete phenomena, he kept a focus on their invariant essences. This resulted in an ontology that sees the world as consisting of objects having properties that are either essential (i.e. defining the category to which the object belongs) or contingent. For example, an object in the category of “humans” has essential

properties, such as having a heart and being able to think, and contingent properties, such as being tired or being located in a particular position.

The next step in the development of the object-based worldview, due to thinkers such as Galileo, Descartes and Newton, was the insight that properties can be measured, i.e. determined in a precise, objective manner as *quantities*. The essential properties remain invariant, but the contingent ones, such as temperature, position or velocity, can vary over the course of time. The corresponding variables define a state space or phase space. Changes in the values of these variables are described by dynamical equations, which express the effect of different forces, such as gravitation, on the objects. Thus, all change can be reduced to the deterministic, time-parameterized trajectory of a collection of invariant objects through their state space (Heylighen, Cilliers, and Gershenson 2007).

This assumption is the foundation of the Newtonian, mechanistic paradigm that has been dominating science at least until the beginning of the twentieth century—when it was put into question by developments such as relativity theory, quantum mechanics, and especially quantum field theory. However, most of the assumptions of this paradigm are still implicitly held by scientists and laypeople alike when they think about how the universe is structured. That is in part because the ontology of objects appears so simple, concrete and intuitive, in part because it is so dominant in the Western, scientific worldview.

Nevertheless, there have always been alternative worldviews that were focusing on processes and relations rather than on invariant objects. These can be found in pre-Socratic and non-Western philosophies, such as those of Heraclitus, Daoism and Buddhism, and in certain ideas proposed by Western philosophers such as Leibniz, Spinoza, Whitehead, Teilhard de Chardin and Bergson. Yet, the success of mechanistic science in explanation, prediction and application was so great that these alternatives were largely ignored. Nevertheless, as our insight into complex, evolving systems advances, the shortcomings of mechanicism and its object ontology become ever more apparent (Heylighen, Cilliers, and Gershenson 2007). To better understand these, we must first formulate the underlying assumptions more explicitly.

The ontology of objects assumes that all phenomena we observe can be reduced to objects and their movement in space. Objects are separate, invariant pieces composed of some inert *substance*, which we call *matter*. Matter may be repositioned, divided or reshaped, but it cannot be destroyed or created. It just stays there, passively waiting until some force may move it into a different spatial arrangement. By assembling elementary pieces of matter, which we call *particles*, into different forms or shapes, different objects can be created. But the objects are intrinsically just as passive, inert or insensitive as the matter that constitutes them. Unless some force acts on them, they remain as they are: rigid, invariant, unchanging.

Objects are distinguished by their location in space. No two objects can occupy the same space: some minimal distance must separate them. Therefore, distinct objects are intrinsically independent: you can move the one without affecting the others. Objects can be identified by their form, i.e. the geometry of the matter that

constitutes them. Form and position determine their properties. Therefore, properties are intrinsic, absolute and objective. They do not depend on the presence or absence of other objects, including those constituting what we call an observer, environment, or context.

Complex objects (“systems”) that are constituted of smaller objects (“parts”) inherit their properties from these parts. There are no “emergent” properties: the whole does not have anything that is more than the sum of its parts. Therefore, to understand the behavior of the whole, it is sufficient to accurately observe the changes in the properties of its parts, i.e. in the variables that describe the system. These changes are driven by external forces obeying natural laws. For example, the fundamental object of neo-Darwinist theory is a population. Its variable properties are the frequencies of different genes. Changes in these frequencies are driven by the force of natural selection, here conceived as adaptation to a given, external environment.

Up to now, the ontology of objects is perfectly consistent. However, if we wish to introduce the observing subject, i.e. the human individual who is perceiving and manipulating these objects, then we are confronted with a paradox. If we assume that this observer can reason and freely decide how to act on these objects, then this observer cannot itself be an inert object. The observing subject must have some form of autonomous agency. This includes a sensitivity to the outside world, some form of motive or goal, and the ability to manipulate and relate to these objects it is observing. Descartes resolved this paradox by postulating that human subjects are constituted out of the very different substance of *mind*, which exists independently of the *matter* that constitutes objects. This ontology is known as *dualism*, because it assumes two separate realms: mind and matter.

While intuitively plausible, from a scientific perspective dualism is a highly unsatisfactory ontology, which creates many more questions and paradoxes than it resolves (Heylighen and Beigi 2018). These largely center on the seemingly unsolvable “mind-body problem”: how can the non-physical mind affect, and be affected by, the inert matter of the body, whose activity is already fully determined by physical forces? Moreover, dualism is incompatible with an evolutionary worldview, which sees the human mind as gradually emerging from much simpler animal, unicellular and inorganic forms of organization.

In earlier work (Heylighen and Beigi 2018), I have proposed a radical resolution to this conundrum, by interpreting both objects and subjects as special cases of the more general category of agents, i.e. phenomena that act on other phenomena. The remainder of this text will spell out some of the implications of this ontology for understanding biological organization and its evolution, by interpreting not only organisms, but their physical components as well as the ecosystems they collectively form, as interrelated agencies. Table 1 provides a first sketch of how agencies differ from objects, a summary that will be elaborated in the rest of the paper.

	Objects	Agencies
<i>Sensitivity</i>	inert	sensitive
<i>Activity</i>	passive	active
<i>Dynamics</i>	mechanical	goal-directed
<i>Shape</i>	rigid	variable
<i>Positioning</i>	localized	localized or distributed
<i>Constituents</i>	material particles	processes
<i>Existence</i>	independent	relational

Table 1: difference between objects and agencies

Animism: the origin of relational agency

The roots of the worldview that I will call relational agency are actually much older than the one of the objecty worldview. The worldview of hunter-gatherers has been characterized as *animism* (Harvey 2014). That means that they see the phenomena around them as animated, i.e. as having agency. These phenomena of course include the animals they hunt, live with, or fear, but also plants, rivers, clouds, and potentially even rocks. To survive and thrive in their local environment, they need to establish good relationships with these other agents. That requires getting to know their habits, capabilities, and preferences, while paying them due respect and trying to keep on good terms with them. They do not see themselves as separate from nature, but as merely one type of agents among many, taking part in a complex network of relationships.

According to the traditional view, animism is a naïve, irrational projection of a human-like mind into physical objects, thus attributing supernatural, “spiritual” properties to these objects. However, a more recent perspective in anthropology (Astor-Aguilera and Harvey 2018; Bird-David 1999; Ingold 2006) interprets animism in a more pragmatic manner, as a practical and adaptive way of understanding the complex ecosystem in which hunter-gatherers live. Supernatural attributions, in this view, are misinterpretations made by old-school anthropologists that were biased by modern religions with their dualist ontology, in which spirits are conceived as disembodied, ghostlike entities separate from the physical objects they may inhabit. The Latin root of the word “spirit” (as well as the “anima” that gave rise to animism), however, simply means “breath”. Breathing is what distinguishes a living, animated body from a dead one. Thus, in the more recent interpretation, animists are merely giving due respect to the “aliveness” of the world that surrounds them.

For hunter-gatherers, that aliveness, animation or agency is not limited to biological organisms. A river for example is an active, self-organizing process that adapts its flow pattern to rainfall, landslides, or mud accretions, while persisting in its goal of discharging its water into a lake or sea. Even a rock may adapt to

circumstances, such as eroding under the effect of frost, becoming slippery in the rain, or providing a substrate for animals to hide in its holes. Thus, a rock too is active: the processes it takes part in are merely much slower than those of animals and plants (Abram 2010). Therefore, it may be seen as animated by some “spirit of the rock”, whose habits and preferences it would be good to know. That would be useful if for example you regularly need to seek shelter underneath that rock, where you may otherwise be surprised by changed conditions, such as burrowing animals or leaking water.

Such detailed knowledge and understanding of natural agencies seem to have helped our hunter-gatherer ancestors to survive and thrive for millions of years before the advent of civilization. In the Paleolithic environment in which humans evolved, animistic thinking appears to have been an adaptive strategy. It would therefore be shortsighted to dismiss animism as merely a pre-scientific misconception. A more sensible approach would be to ask why we eventually felt the need to shift from animism to our present object-based worldview.

The origin of the object ontology

The anthropologist and philosopher David Abram has argued that this shift followed the adoption of alphabetic writing—first by the ancient Hebrews, then by the Greeks (Abram 2012). Expressing ideas in the form of abstract symbols registered in an enduring document separates these ideas from the concrete, variable phenomena to which they refer. The shape of letters and words moreover has no relation with the shapes of these phenomena—another disconnect. Furthermore, the words registered on a page are intended to be read at a different time and place, i.e. in a different context, than the one in which they were written. Therefore, their meaning, i.e. the phenomena they refer to, is supposed to be invariant across changes in context (Heylighen and Dewaele 2002). That means that it stands on its own, independently of other phenomena. Thus, according to Abram (2012), written words became the prototype of the eternal, absolute “forms” postulated by Plato to be the fundamental constituents of reality. These in turn led Aristotle to conceive of objects as the material embodiments of these invariant and independent essences.

It is this context-independence of written descriptions that allowed knowledge to be recorded for the long-term and to be disseminated widely without losing its meaning (Heylighen 1999). Thus, for example, ancient Greek treatises about astronomy, geometry and philosophy could still be read and understood two thousand years later in the very different context of the European Renaissance, providing an impetus for the scientific revolution. The further development of mathematical formalisms, such as Cartesian coordinates, calculus and algebra, helped scientists to describe the properties of objects increasingly precisely and objectively, making abstraction of the observing subject and the context of observation. The resulting formalization allowed knowledge to be expressed in the form of universal “laws of

nature” that are independent of space, time or observer, thus forming the foundation of the classical scientific worldview (Heylighen 2012).

Such precise, context-independent description moreover facilitated the development and spread of technology, because it allowed engineers to write down a full specification of the components and mechanisms of the objects they designed. Yet, even before formal descriptions, the object-based way of thinking had been stimulated by our use of an ever-larger array of artifacts, such as furniture, pottery or tools. Indeed, unlike natural systems, artifacts are designed to be mere instruments, i.e. passive objects that can be dependably manipulated by their human users—without deploying any agency of their own. You would not want your table to move to a different room because it prefers the view there! Neither would you want your hammer to decide it is better to swing sideways when you want it to swing downwards...

As people spent more and more time inside artificial environments, such as houses, cities and offices, their interactions with the environment were increasingly restricted to the perception and handling of such passive objects. This stands in sharp contrast with the natural environment inhabited by hunter-gatherers and early farmers, where most phenomena, such as plants, animals, forests, rivers and clouds, are intrinsically animated. Those, such as pebbles or sticks, that do behave like inert objects appear to be the exception rather than the rule. In our present environment, on the other hand, we are surrounded by passive objects, and it are non-human agencies that have become the exception. As a result, we have great difficulty conceiving and experiencing nature as fundamentally active and relational.

Thus, it seems as if the development of civilization, and especially of the symbolic systems supporting science and technology, has alienated humans not just from their natural environment, but from the animistic awareness that makes us intuitively understand and feel part of that living nature (Charlton 2007). Abram and others (Abram 2010; Harvey 2019) have argued that our actions have severely upset the ecological balance—exhausting resources, polluting rivers, and driving many species to extinction—in large part because we see nature as a mere collection of objects to be manipulated, rather than as a network of agencies in which we are intimately involved by a variety of relationships.

A renaissance of relational agency?

Is there a way to recover the relational understanding needed for a deep ecological consciousness, yet without losing the benefits of objective, scientific description? That would require a conceptual framework and formal language able to accurately represent agencies and their relations, without reducing them to independent objects. The core message of this paper is that the elements of such a description are at hand: they merely need to be further clarified, elaborated and integrated. Let us illustrate this development with some recent approaches.

One example of a presently popular formalism is *Agent-Based Modeling* (Bonabeau 2002; Macal and North 2010). This pragmatic method is used to simulate the often-complex interactions between autonomous agents, and thus understand the emerging patterns of their joint evolution. Agents here are conceived as any entities that act on each other and the environment they share. These can be organisms, cells, molecules, firms, robots, or people. Computer simulation, unlike written language, can represent actual processes, showing us in real time how actions result in further actions. Such *multi-agent simulations* (Ferber, Gutknecht, and Michel 2004; Macal and North 2010) are commonly used to illustrate a theoretical perspective known as *Complex Adaptive Systems*. Typical examples of such systems are ecosystems, markets, and societies. Their dynamics combines individual variation and selection, adaptation, competition, cooperation, and collective self-organization (Hartvigsen, Kinzig, and Peterson 1998; Holland 2012; Miller and Page 2007).

A related conceptual framework, *Actor-Network Theory*, originates in sociology (Latour 1996). It describes societal evolutions, such as the development of nuclear energy, as resulting from the interaction between very diverse “actors”. These include human individuals, organizations, institutions, technologies and even physical resources such as uranium, while forming a network of mutual dependencies. In the humanities, a similar perspective has inspired a new philosophical stance known as *Posthumanism*. Posthumanists argue that, when reasoning and deciding about our common future, we, humans, should take into account our relations with various non-human agents, such as animals, plants, robots, or rivers (Barad 2003; Braidotti 2013).

This brief sample of recently developed approaches illustrates the potential of relational agency to become not just a philosophical perspective, but a scientific method for investigating complex, evolving systems. Nevertheless, these existing approaches remain as yet fragmentary, disparate, and incoherent, originating from widely divergent viewpoints, disciplines and theoretical frameworks. The same applies to the different “extensions” of evolutionary theory that we reviewed, to domains such as ecology, development and molecular biology. In order to unify these different approaches, we need a conceptual framework and language that can deal with relations and agencies at all levels, from molecules to organisms, ecosystems and societies. In the remainder of this paper, I will propose a number of concepts and modeling methods that may serve as a foundation for such an integrated theory of relational agency.

The ontology of actions

The ontology of objects assumes that there are elementary objects, called “particles”, out of which all more complex objects—and therefore the whole of reality—are constituted. Similarly, the ontology of relational agency assumes that there are elementary processes, which I will call *actions* or *reactions*, that form the basic constituents of reality (Heylighen 2011; Heylighen and Beigi 2018; Turchin 1993).

A rationale for the primacy of processes over matter can be found in quantum field theory (Bickhard 2011; Kuhlmann 2000). Quantum mechanics already made clear that observing some phenomenon, such as the position of a particle, is an action that necessarily affects the phenomenon being observed: no observation without interaction. Moreover, in general the result of that observation is indeterminate before the observation is made: the action of observing in a real sense *creates* the property being observed, through a process known as the collapse of the wave function (Heylighen 2019; Tumulka 2006). For example, before the observation a particle such as an electron typically does not have a precise position in space, but immediately afterwards it does. More generally, quantum mechanics tells us that microscopic objects such as particles in general do not have objective, determinate properties, but that such properties are (temporarily) generated through interaction (Barad 2003).

Quantum field theory adds that the objects (particles) themselves do not have a permanent existence, but that they too can be created or destroyed through interactions, such as nuclear reactions. Particles can even be generated by fluctuations of the vacuum (i.e. out of nothing), albeit that in this case their existence is so transient that they are called “virtual” (Milonni 2013).

At a larger scale, the molecules that constitute living organisms are similarly ephemeral, being both produced and broken down 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. Similar creation-destruction processes can be found at the level of ecosystems, where relations of predation, symbiosis, and reproduction between organisms join with meteorological and geological forces to produce a constantly changing landscape of forests and rivers, mountains and meadows. Even the planet is in a constant flux, as magma moves up, down and sideways under its mantle, moving continents and opening up volcano-studded rifts. And the sun and stars that irradiate it are merely boiling cauldrons of nuclear reactions generating new elements in their core, while releasing immense amounts of energy.

Against Parmenides, we may conclude that at all levels, from particles to stars, change (“Becoming”) is more fundamental than permanence (“Being”) (Prigogine 1980), and that the stability of macroscopic objects is only apparent, ready to be dissolved when we more closely observe the microscopic processes that constitute them. That is why contemporary science needs a formal language that treats change as primary, and stability as derivative.

Reactions, actions and agencies

Possibly the simplest way to represent an elementary process may be called a *condition-action rule* (Heylighen and Beigi 2018). It notes that given some condition *X*, an action will take place that produces the new condition *Y*. This can be interpreted

as a transformation from the situation X to the new situation Y , or as a causation linking the cause X to its effect Y . It can be expressed more briefly as “*if X , then Y* ”, or in symbols:

$$X \rightarrow Y$$

Conditions here are any states of affairs whose presence can be distinguished from their absence. Examples of conditions are the presence of various types or species of particles, chemical compounds, living organisms, habitats, geological features, or meteorological circumstances. An example could be “*if dark clouds are gathering, then it will rain*”.

These conditions can in general be decomposed into conjunctions of more elementary conditions, such as the joint presence of two or more circumstances, particles, chemicals or biological species. For example, “*if dark clouds are gathering and people are in the streets, then people will seek shelter and it will rain*”. Adopting the notation used for reactions in physics and chemistry, I will denote conjunctions (“*and*”) by the “+” operator:

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

In this notation, the condition-action rules can be interpreted as *reactions*, in the sense that the conditions $\{a, b, \dots\}$ on the left-hand side of the arrow “react” with each other in order to produce the novel conditions $\{e, f, \dots\}$ on the right-hand side of the arrow. Such a reaction represents an elementary process that takes the elements on the left as its inputs and transforms them into the elements on the right as its outputs (Heylighen, Beigi, and Veloz 2015).

Agencies (A) can now 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 agency of the conversion. The *reaction* between A , X and Y can therefore be reinterpreted as an *action* performed by the *agency* A on condition X in order to produce condition Y . This can be represented in shorter notation as:

$$A: X \rightarrow Y$$

Note that while an agency remains invariant during the reactions it catalyzes, there in general exist other reactions that destroy (consume) or create (produce) that agency.

Thus, while agencies can play the role reserved in the old ontology for objects or forces, they are neither inert nor invariant.

An agency will in general catalyze several reactions. This means that it will react to different conditions by different actions so as to produce different new conditions. For example:

$$A: X \rightarrow Y, Y \rightarrow Z, U \rightarrow Z$$

This set of actions triggered by A can be interpreted as a *dynamical system* that maps initial states (X, Y, U) onto subsequent states (Y, Z, Z) (Heylighen 2022; Sternberg 2010). Such a dynamical system typically has one or more *attractors*. These are (sets of) states to which the dynamics converges. That means that the processes catalyzed by the agency A will lead into the attractor, but that A does not provide any way out: once in the attractor, further actions can only produce states that are also in the attractor. An attractor is surrounded by a *basin of attraction*. This contains the states that converge to the attractor, but that are not part of the attractor itself. In the above example, Z is an attractor, while X, Y and U are in its basin. That means that processes starting from these conditions necessarily end up in Z , without possibility of returning.

Such an attractor-centered dynamics can be interpreted as a model of *goal-directedness* or *teleonomy* (Heylighen 2022; Heylighen and Beigi 2018): the actions executed by agency A are directed towards the goal Z . That means that Z constitutes the shared final state or “end” of processes starting from the different initial states X, Y and U , a feature called *equifinality* (Lyman 2004) or *plasticity*. Moreover, perturbations, i.e. processes not controlled by A , that make a state deviate from its goal-directed trajectory, will be counteracted by A so that they still end up in the attractor, at least as long as the deviation remains within the basin of that attractor. This is a characteristic of goal-directed systems that has been called *persistence* or *regulation* through negative feedback (Heylighen 2022).

Thus, an agency can be seen as the enabler of a bundle of actions that act on the present situation so as to drive it towards a restricted set of attractor states. These define the implicit goal, “preference” or end of the agency. While the reasoning that led to this definition is very abstract and general, the definition fits in with our intuitive sense of agency as the capability for goal-directed action. For example, bacteria exhibit chemotaxis, which means that they seek food by swimming up gradients of food molecules towards the source (Sourjik and Wingreen 2012). The source (region with the highest concentration of food) functions here as the attractor of the chemotaxis dynamics. This dynamics is governed by a simple condition-action rule obeyed by the bacterium: *if* the concentration of food molecules does not increase, *then* change direction of movement. The net result is that the bacterium will be persistently moving in a direction where food concentration increases, until it reaches the highest concentration.

In this example, the agency is located within a spatially delimited system: the bacterial cell. We may call such a locally bounded agency an “agent”. For example, the bacterium is the agent of its food-directed movement. An agent is in that respect similar to an object. However, there are also agencies that lack such a clear localization. Examples are gravity, an ecosystem, a social institution, or the weather. For example, gravity plays the role of agency in the following reaction (Heylighen 1999):

unsupported heavy object + gravity → falling object + gravity

The attractor of such gravity-induced dynamics is the lowest position the object can reach until it is supported again. We may say that gravity as an agency “prefers” lowering the positions of heavy objects.

Another example of an agency that is difficult to delimit is a thunderstorm. Its action consists in producing rain, wind and lightning, while its “goal” is to discharge the water it carries onto the earth. In earlier times, this agency may have been conceived as a god of thunder and rain, such as Thor or Zeus. In present times, we try to understand it by means of very complicated meteorological models. But in any case, the impact of this agency on its surroundings is great, and therefore it is important to be able to anticipate and adapt to its actions. This requires an understanding of that agency’s typical behavior. That knowledge includes rules such as that people walking outside during a thunderstorm will get wet, that tall structures may be hit by lightning, and that rivers may flood. An object-based ontology is totally unhelpful in this respect, because a thunderstorm is not an object in any practical sense of the word, while it is useless to try to reduce it to the more object-like molecules of water and air that it contains.

Agents and challenges

From the point of view of an individual agent (or more general agency), the environment consists of conditions to which it may or may not react by performing an action. If the present condition is not a precondition for one of the reactions in the repertoire of condition-action rules that characterize the agent, then that agent will simply ignore that condition. For example, the presence of nitrogen in the air will be ignored by animals, because free nitrogen does not react with any of their metabolic processes. Therefore, they are insensitive to the amount of nitrogen present, and would not notice if that nitrogen is replaced by a different non-reactive gas, such as argon. On the other hand, they are highly sensitive to the amount of oxygen in the air, because they need it for their respiration, and to produce energy. Still, there are nitrogen-sensitive agencies, namely bacteria that can absorb nitrogen from the air and transform it into ammonia or nitrates, which in turn can be used by plants as fertilizer.

Some conditions, such as the presence of oxygen, do trigger reactions from the agent, such as breathing. I will call such activating conditions *challenges*, because they incite or challenge the agent to act, i.e. to change something about the situation (Heylighen 2012; 2014). For a goal-directed agent, a challenge may be positive or negative. It is positive when it helps the agent to get closer to its goal. It is negative, when it makes it more difficult for the agent to reach its goal. For example, for a bacterium the ultimate goal at which its actions are directed is survival and multiplication. The presence of food is a positive challenge, because consuming that food will help the bacterium achieve that goal. The presence of a toxin, on the other hand, is a negative challenge, because it makes it less likely for the bacterium to survive. Therefore, the bacterium will act so as to approach and ingest food and to evade toxins.

A positive challenge may be called an *affordance*, because it provides the agent with an opportunity to achieve some benefit, i.e. move closer to its goal. A negative challenge may be called a *disturbance*, i.e. something that pushes the agent away from its goal-directed course of action. A challenge can also be neutral, in the sense that the action it elicits is neither one of approach nor one of avoidance, but e.g. of exploration. In that case, I may call it a *diversion*: it incites a deviation from the present course of action that brings the agent neither closer to, nor farther away from, its goal.

That course of action is in general unpredictable because the environment will constantly throw up challenges that demand some novel action from the agent, so as to remain on course to the goal (Heylighen 2012). The reason is that the environment consists of other agencies. These similarly react to challenges by actions that produce new conditions. These in turn may challenge one or more agent to act. Thus, we may say that *challenges propagate*: the reaction to a challenge by one agent will typically create a challenge for one or more other agents, potentially setting in motion a cascade of actions triggering further actions (Heylighen 2014). That is because agents share a common environment, which functions as a medium interconnecting the different actions and agencies, thus providing a channel for (implicit) communication and coordination of actions (Heylighen 2016). Let us investigate the resulting network of relationships.

Relations between actions and agencies

Unlike objects, the reactions that form the elements of our new ontology are intrinsically coupled. That is because the conditions that constitute the inputs and outputs of a reaction always must occur in some other reaction as well. This is the principle of “the difference that makes a difference” (Bateson 2000): for some condition to be distinguishable it must incite a reaction that itself leads to a distinguishable condition. A condition that does not produce any distinguishable effect simply cannot be observed, neither directly nor indirectly. Therefore, we may

as well assume that it does not exist. For example, it is in principle possible to reintroduce inert objects in the action ontology by means of the reaction $a \rightarrow a$ (the object a is always followed by the object a , without any change). But if we want to observe whether object a is present, we will need at least a reaction of the form: $a + \text{observation} \rightarrow a + \text{observation result}$. For example, the observation may consist in shining light on the object, while the observation result consists in the detection of light reflected back by the object.

Therefore, different reactions necessarily share certain conditions, in the sense that the inputs or outputs of one reaction must also be outputs or inputs of some other reaction. That means that reactions are mutually dependent. If, for example, one reaction *produces* the condition X , while another reaction *is triggered by* the condition X , then the first one will automatically be followed by the second one. Yet, relations can be more complicated. For example, one reaction may be triggered by $a + b$, while another is triggered by $a + c$. If now some third reaction produces a , then, depending on the presence or absence of b and c , none, one or both of these initial reactions may be triggered. Thus, reactions form a complex network of interdependencies (Veloz and Razeto-Barry 2017).

Some of these relationships will be synergetic, in the sense that two or more agencies or reactions together can produce more of the conditions or resources they all need to continue functioning than each of them on its own. Others will be characterized by conflict or friction, in the sense that the activity of the one will impede the continued activity of the other(s). An agency surrounded by synergetic agencies will maintain more easily than one surrounded by agencies that have a relation of friction with it. Therefore, natural selection will tend to favor agencies profiting from synergies and eliminate agencies suffering from frictions. Thus, there will be a general trend for evolution to promote synergetic relationships between agents, while weakening relationships characterized by conflict or friction. However, in a complex network of relationships involving several agencies, as we find in ecosystems or biological networks, it is not always obvious what precisely constitutes synergy. Let us illustrate such mutual dependencies with a simple example of a self-contained ecosystem, and then elaborate it by introducing variation and selection, so as to develop a relational perspective on evolution.

Self-maintaining systems

Consider a special type of aquarium that does not exchange any matter with the outside world. Such a hermetically sealed, transparent bowl, called an “ecosphere”, contains air, seawater, shrimps, algae and bacteria (Schilthuizen 2008). The organisms in such a bowl can survive for years without requiring any outside intervention. That is because their activities form a self-maintaining whole that can be summarized by the following reactions:

→ light

shrimps + algae + O₂ → shrimps + waste + CO₂ + heat

bacteria + waste + O₂ → bacteria + nutrients + CO₂ + heat

algae + light + CO₂ + nutrients → 2 algae + O₂

heat →

These reactions describe how the shrimps consume algae and oxygen, while producing waste and carbon dioxide. The bacteria transform the waste into nutrients, such as nitrates and phosphates, that act as fertilizer for the algae. These nutrients together with carbon dioxide allow the algae to grow, while using light for photosynthesis and producing oxygen as a byproduct. The additional algae feed the shrimps, so that these can survive, and sustain the initial reaction.

This network of reactions includes three agencies: shrimps, algae and bacteria. Their actions are perfectly complementary: each produces what the others consume, and consumes what the others produce. The only resource that needs to come from outside the system to keep it going is light. Its only output is heat. That is because such a self-maintaining system, even when it recycles most of the products it uses, must be thermodynamically open in order not to degrade to its maximum entropy state in which everything decays.

It may seem strange that the amount of resources produced and consumed would remain perfectly balanced, so that the system maintains itself over the years. The explanation is that the different agencies will adjust their production and consumption of resources until they are mutually adapted. Suppose that initially there are too many algae for the amount of nutrients available. More algae means more food for the shrimps and therefore more waste and thus nutrients produced. Thus, the amount of algae will diminish by increased consumption, while the amount of nutrients they need will increase, until both are in balance. Suppose now that there are not enough algae initially. This may lead to the dying off of some of the shrimps. The bacteria will decompose the dead shrimps into nutrients. Fewer shrimps means less consumption of algae. This together with the additional nutrients will allow the algae to multiply, until their number is in balance with that of the other components of the system.

More generally, every fluctuation up or down of one of the resources that make up this system will elicit a feedback that is negative, i.e. that moves in the opposite direction of the initial fluctuation. The reason is that in such a self-maintaining network of reactions the reduced availability of some resource will normally lead to a reduced consumption of that resource, while production of that resource does not immediately decrease, so that the net availability increases again. A similar negative feedback will reduce increased availability. Thus, resource concentrations, while they can vary, generally do not deviate too far from their equilibrium values.

Such a closed ecosystem illustrates a very important concept in reaction networks: a (*chemical*) *organization*. This concept was introduced by Peter Dittrich,

thus founding an approach known as Chemical Organization Theory (Dittrich and Fenizio 2007; Heylighen, Beigi, and Veloz 2015; Veloz et al. 2022). In this theory, an organization is defined as a network of reactions and resources (also called “molecules” or “species”) that is *closed* and *self-maintaining*. *Closed* here means that no resources are produced by the reactions that are not already part of the network. *Self-maintaining* means that all resources that are consumed by some reactions are produced at least as much by other reactions, so that their total amount does not go to zero. In other words, the resources in a chemical organization are fully recycled and thus remain perpetually available.

Interestingly, it can be shown that the *attractors* of the dynamical system defined by the network of reactions are all chemical organizations (Peter and Dittrich 2011). That means that the system tends to spontaneously settle in one of these organizations, and that once there, it will remain there. In other words, such self-maintaining, closed networks tend to self-organize. Self-organization, as observed long ago by the cybernetician Ashby, can be conceptualized as the mutual adaptation or alignment of the different subsystems within an overall system (Ashby 1962; Heylighen 2001). Let us illustrate the principle by imagining how one could create a new self-contained ecosystem from scratch.

Introducing variation and selection

Instead of buying a commercially available “ecosphere” aquarium, you might prefer to start one of your own. A simple method is to fill a big glass jar with a variety of organisms and materials collected in ponds, such as soil, pond water, water plants, small crustaceans, snails, worms and insect larvae, and perhaps even some tadpoles or small fish. You then close the jar so that the water cannot evaporate, leaving some air at the top, leave it in a place with sufficient light, and wait. Typically, many of the organisms will die, because they lack the right conditions necessary for their survival, or because they do not find enough resources, such as food organisms, to sustain a thriving population. If all the members of a species have died, so that no new ones can be produced, then that species has been removed forever from the closed ecosystem. This is *natural selection* at work on the micro-scale of an ecosphere: that species was not adapted to its local environment.

Yet, some species are likely to survive. Initially, their populations may fluctuate wildly, but then overproduction of one species (e.g. overgrowth of algae profiting from the nutrients released by the decomposition of dead organisms) will tend to be suppressed by the corresponding population growth of another species that consumes them (e.g. snails eating algae). As this negative feedback sets in, fluctuations decrease, and the system reaches some degree of stability. Note that the resulting self-sustaining ecosystem may not be very interesting to look at: larger and more active organisms, such as fish, are less likely to find the necessary resources to survive in such a small volume. But, speaking from personal experience, some of the

plants, microorganisms and invertebrates will survive for years without requiring any outside intervention. These species are *fit*, i.e. adapted to the closed environment they collectively constitute.

In terms of Chemical Organization Theory (COT), this means that they are produced as least as much as they are consumed. When the rate of the different reactions in a reaction network can be estimated, the network can be modeled as a dynamical system. That allows you too simulate the evolution of the ecosystem, by calculating how the concentrations of the different resources or species go up and down in the course of time. Such simulation will typically end in an attractor regime, where some of the concentrations have gone down to zero (elimination), while others remain positive (survival), though they may continue to fluctuate (Veloz et al. 2022; Heylighen, Beigi, and Veloz 2015). Even without such a quantitative simulation, COT often allows you to determine algebraically which of the species will necessarily be eliminated—e.g. because the reactions producing it are insufficient to compensate for the reactions consuming it—but also which assemblies of species are capable of forming a self-maintaining organization.

This shows that fitness can in principle be defined in a purely relational manner: as the ratio of production to consumption in a network of reactions connecting different resources and agencies. That implies that the fitness of an agency is context-dependent: it can be large or small depending on the presence or concentration of other agencies and resources. Moreover, the definition of fitness in a given context is ultimately circular or “bootstrapping”: the fitness, and therefore presence, of these other agencies in turn depends on the presence or fitness of the first agency. That is because they are all part of a closed organization of mutually feeding reactions. For example, in the above example of the ecosphere, the fitness of the shrimps depends on the fitness of the bacteria, because the shrimps need the bacteria to turn a potentially toxic buildup of waste into nutrients that will feed the algae that they consume. Vice-versa, the bacteria need the shrimps to eat and digest the algae, turning them into waste that they can feed on.

Such mutual dependencies are usually conceived as cycles, in which some resource x is turned by some agency into y , after which one or more other agencies turn it back into x . But the reaction network formalism can handle much more complex mutual dependencies, where resources are produced again via a complex, branching network of reactions that turn combinations of resources into other combinations.

Let us further expand the example of the self-maintaining ecosystem. At the largest scale, the planet Earth is a closed ecosystem, which only receives sunlight as an input, while radiating heat into space. The number of species and resources making up this ecosystem is immensely larger than the one in our closed jar. Yet, the network of reactions describing their mutual dependencies is not fundamentally different. Both are characterized by processes of mutual adaptation and overall self-organization that generate a self-maintaining, closed whole. At the coarse level of the biosphere, we even find a similar balance between oxygen-producing plants, carbon dioxide-

producing animals and nutrient-producing decomposers—although here we also need to take into account processes of recycling in the atmosphere, oceans and lithosphere (e.g. carbon dioxide being trapped in sediments and chalk formations). This self-maintaining system at the planetary scale is known as Gaia (Rubin, Veloz, and Maldonado 2021).

Perhaps the most striking difference at this scale is that we can expect frequent processes of *variation*: agencies undergoing a transformation that changes the way they react to certain conditions. For example, imagine that some species of bacteria undergoes a mutation that makes it resistant to a common antibiotic. This species was being produced by a reaction of the form:

$$\text{bacteria} + \text{host} \rightarrow 2 \text{ bacteria} + \text{ill_host},$$

while being consumed, and thus held in check, by a reaction of the form:

$$\text{bacteria} + \text{ill_host} + \text{antibiotic} \rightarrow \text{host}$$

The mutated form, *bacteria**, however, is no longer consumed by the latter reaction, and thus may spread exponentially via the first reaction—unless some new reaction kicks in to consume it, such as:

$$\text{bacteria}^* + \text{host} \rightarrow \text{host} + \text{antibodies} + \text{bacteria}^* \rightarrow \text{host}$$

Variation creates new types or species of agencies, which may or may not fit in with the existing ecosystem. Unfit variations are simply eliminated. However, fit variations change the ecosystem itself, by initiating reactions that produce new challenges for other agencies. These may shift the balance between consumption and production of resources and agencies, potentially leading to the extinction of hitherto fit species. On the other hand, these changes may carve out a stable niche for the newly evolved species. Thus, a new variation resets, in a minor or—more rarely—major way, the dynamical system defined by the reaction network, inciting it to converge to a different attractor regime.

Cells and genes as agencies

After zooming out to the largest scale of the planetary ecosystem, I will now zoom in to the smallest biological scale, the one of a single-celled organism, such as a bacterium. In this I am following the old dictum “As above, so below” (ascribed to the mythical sage Hermes Trismegistus), which suggests a fundamental similarity between macrocosm and microcosm.

Indeed, from a relational agency perspective, a living cell is not essentially different from an ecosphere. Both are to some degree transparent containers, filled

with a diversity of interacting “molecules” or “species” that are continually being consumed and produced, together forming a self-maintaining, closed network of reactions. Thus, the overall system is *autopoietic*, i.e. self-producing (Heylighen, Beigi, and Veloz 2015; Mingers 1994; Razeto-Barry 2012). The complex, branching cycles of consumption and production reactions constitute the cell’s *metabolism*. While we as yet do not know all the reactions in a bacterial cell’s metabolism, those parts we do know, such as *E. coli*’s sugar metabolism, are easily modeled as chemical organizations (Centler et al. 2007).

In the cell, the role of the resources is played by “passive” molecules, such as glucose, ATP and oxygen, which are consumed and produced by reactions in order to harness energy or build components. The role of the agencies is played by the enzymes, which catalyze and thus enable most of these reactions. Yet, the enzymes themselves are the product of more complex gene-expressing processes, which read a coding sequence of DNA and translate it into the right enzyme. As long as the cell does not divide, the DNA itself is invariant, i.e. neither produced nor consumed. That allows it to function as a dependable memory for the cell, storing the “knowledge” about which actions (enzyme-catalyzed reactions) to perform in order to deal with various conditions (presence or absence of particular molecules in the cell).


For example, the antibiotic-resistant bacterium may have acquired a piece of DNA that codes for an antibiotic-neutralizing enzyme. The entry of the antibiotic into the cell then functions as a triggering condition to express that stretch of DNA into the corresponding enzyme. That enables a reaction that consumes the antibiotic, transforming it into a molecule that is no longer toxic for the cell.

From this perspective, the genes in a cell are functionally similar to the different agencies in a self-maintaining ecosystem. The reactions they trigger must be coordinated or mutually adapted, so that the one produces what the others consume, and the system as a whole remains in balance—while being quick to adjust to fluctuations in the concentrations of the different molecular “species”, or to the entry of food molecules that must be consumed or toxins that must be neutralized. That means that for a cell to survive, its genes must efficiently cooperate, forming a synergetic whole where everything needed by one process is produced by one or more other processes, and vice versa. Note that these processes generally include reactions that transport resources into the cell, and waste out of the cell. Therefore, the system is not closed in the thermodynamic sense, only in the organizational sense.

From a “selfish gene” perspective (Dawkins 2006), it is not obvious to explain how all these agencies would have learned to cooperate. From a relational agency perspective, on the other hand, inputs and outputs of the different agencies will eventually adjust to each other, albeit that some of the agencies may be eliminated in the process—as illustrated by the ecosystem in a jar where some of the organisms do not survive. The remaining ones can be viewed as cooperative, since it is through their on-going contributions that the overall, autopoietic system is maintained. Note that this conclusion is not in contradiction with the selfish gene perspective, which sees genes shared by an organism as being “in the same boat” (Dawkins 2006)

(surviving or being eliminated together), and therefore forced by natural selection to work together effectively. What the reaction network model adds is a mechanism clarifying what “working together” precisely means, and how such an arrangement can evolve.

What the cell example further adds to the one of a self-maintaining ecosystem is that the cellular organization produces its own boundary, the cell membrane. This membrane is semi-permeable: it allows “food” molecules to enter, and “waste” molecules to exit, but otherwise protects the self-producing metabolism from external molecules that may interfere with its reactions. That makes the autopoietic organization more robust or resilient (Veloz et al. 2022), and therefore less dependent on outside conditions, i.e. more autonomous. Its implicit goal is self-maintenance (and eventual replication). Therefore, we can view the cell as a whole as an agency in its own right, a “super-agency” consisting of genes and enzymes as “sub-agencies” that control its internal metabolism. This agency acts on its environment by transforming certain available inputs into outputs that change the environmental situation. If we use a unicellular alga as an example, we are back to the reaction that started this series of examples:



This shows that the relational agency ontology, with its formalization in terms of reaction networks and Chemical Organization Theory, is applicable at all levels of life, from molecules to the planetary ecosystem.

An extended evolutionary synthesis?

I have presented a first sketch of a new framework that can help us to understand complex, interacting processes and the systems they produce. This framework functions both on the conceptual level and on the level of formal descriptions that can be investigated mathematically and through computer simulation. Because these descriptions start from the very general concepts of conditions and actions, they do not a priori distinguish between microscopic processes (such as reactions between molecules) and macroscopic processes (such as interactions between biological species and their environments). That makes them in principle very widely applicable, and thus potentially able to unify disparate fields (Heylighen, Beigi, and Veloz 2015; Veloz and Razeto-Barry 2017), both inside and outside biology. Let us here examine whether such a framework could support an extended evolutionary synthesis, whose divergent strands I sketched in the introduction.

First, the relational agency framework embraces *self-organization* from the very beginning: following the lead of Ashby, it sees the self-organization of a complex system as the mutual adaptation of the components of that system so as to develop a self-maintaining network of relationships. Thus, self-organization provides

the system with “order for free”, as Kauffman has called it, i.e. without need for selection by the environment to produce such order (Kauffman 1993). Yet, we should note that self-organization does imply a form of *internal selection*, in the sense that agencies or resources that do not manage to adapt to the rest of the emerging organization are eliminated. Moreover, of all the possible dynamical arrangements of reactions only the self-maintaining organizations survive for the long term. The resulting mutually dependent reaction networks form the main focus of *systems biology*, an approach that examines how processes within an organism work together to keep the metabolism going.

From the perspective of relational agency, evolution is always *co-evolution*: there is no a priori distinction between the system that is evolving and the environment to which the system has to adapt. Indeed, an environment consists itself of agencies that are trying to adapt to the environment created by all the others. Thus, adaptation in the one will in general challenge the others to adapt as well, potentially triggering a cascade of propagating challenges and concomitant adaptations. These processes of adaptation are in general not passive, but active: an agency intervenes in its environment by consuming inputs and producing outputs, thus changing the conditions for itself and others. These actions are not just mechanical responses to external forces, but directed at the internal goal of self-maintenance. The reason is that agencies whose actions were not effective in safeguarding the continuation of their organization have simply been eliminated by natural selection. Therefore, the remaining agencies are goal-directed or *teleonomic*: their repertoire of condition-action rules is such that a wide variety of initial conditions is made to converge to the same attractor regime of self-maintenance (Heylighen 2022; Mossio and Bich 2017). The short-term effect of these actions is to neutralize perturbations that make the agency deviate from its autopoietic regime. The long-term effect is to reshape the environment so that its conditions make self-maintenance easier—in other words, to *construct a niche* for the organization.

Another long-term effect is to establish dependable relationships with the other agencies in the co-evolving network. These relations are ideally *synergetic*: the agencies mobilize more resources for their continuing self-maintenance together, i.e. in interaction, than they would have done when acting on their own. This synergy results from self-maintenance at the level of the collective organization, where some agencies produce what others consume, and vice-versa. Thus, the agencies live in *symbiosis*, being to some degree dependent on each other’s activities. This symbiosis is not necessarily mutualistic, but can encompass parasite-host or predator-prey type relationships, where one species grows at the expense of another species—like in the example of shrimp consuming algae. The synergy in a self-maintaining reaction network does not require consumption to be immediately repaid by production. The contribution of the consumer (e.g. shrimp) to the producer (e.g. algae) can be indirect, via a branching pathway that involves several intervening agencies (e.g. bacteria) and their products (e.g. nutrients).

When the system of relationships becomes so strong that an agency no longer can afford to survive without it, the self-maintaining network starts to behave like an individual agency rather than a group of collaborating agencies. In this way, agencies become integrated into a super-agency. If this agency is localized within a clear boundary, it forms an agent. When such a super-agent is constituted of organisms that belong to different species, the process is called *symbiogenesis*: the generation of a new type of organism by the integration of symbiotically living organisms. When they belong to the same type or species, the process is called an *evolutionary transition*. For example, cells can develop more synergetic relations by specializing in complementary functions, such as somatic versus reproductive, thus forming a differentiated, multicellular organism. Since self-maintenance can succeed or fail as well for the super-agent as for its constituent agencies, there is now selection at both levels. Thus, there is a continuing pressure for *evolution at multiple levels* to produce condition-action rules that make the respective agencies fitter.

I have not as yet touched on the more complicated issue of ontogenetic *development*, i.e. the process during which the cells of a growing embryo differentiate into various tissues and organs. Kauffman's influential work on self-organization and selection in evolution attempted to understand the role of genetic regulatory networks in guiding this differentiation (Kauffman 1993). He modeled these networks as random Boolean networks, because that made it easy to simulate their self-organizing dynamics. This dynamics has multiple attractors. These can be interpreted as different cell types characterized by different combinations of genes being "on" (expressed, active) or "off" (silent). Development can then be understood as a process in which different cells spread out across the different valleys of their epigenetic landscape (Baedke 2013), so as to end up in their respective attractors.

It seems likely that random Boolean networks can be expressed in the same language of reaction networks that we have been using to describe elementary processes. This language is actually richer, because it can deal not only with Boolean (binary) variables, but also with concentrations expressed as real numbers. Therefore, it may well provide a more realistic model of the networks that are formed by genes, functioning as agencies, that regulate reactions and each other's activities via enzymes. There is no reason to assume that the attractor landscape formed by the different self-maintaining "chemical organizations" in such a reaction network would be any less rich than the one characterizing a random Boolean network. Thus, it seems as if the relational agency paradigm holds the potential to model the self-organizing evolution of not just cells and ecosystems, but of multicellular organisms and their development.

Conclusion

While the theory of evolution is fundamentally about change, it is remarkable how much the standard version of it, neo-Darwinism, is focused on static entities: genes

and environments. In that respect, it is not different from other scientific theories of change, such as Newtonian mechanics. That is because all these theories ultimately derive from an ontology dating back to Aristotle, which sees invariant objects as the most fundamental constituents of reality (Walsh 2018). Change, then, is nothing more than the movement of these objects through some abstract space defined by their variable properties, which is driven by external forces, such as gravity or natural selection. The advantage of that approach is that objects, their properties and the forces that drive them are easy to express in a precise, context-independent manner, allowing formal modeling.

Yet, there is an alternative ontology, of which the earliest incarnation is known as animism, but which I have here called *relational agency*. It sees the basic constituents of reality as elementary processes that define a network of interdependent agencies. Until recently, such process philosophy was hard to express in a precise manner. Therefore, it had little impact on scientific theory, which is still dominated by the object ontology. At present, however, approaches such as agent-based modeling and reaction networks have provided us with the beginnings of a formal language and method for computer simulation. This puts us in a position where we can start to develop an integrated theory of evolution that can deal with relational and agential processes.

An elementary reaction can be modeled as a condition-action rule, stating how a particular conjunction of conditions is transformed into some new conjunction of conditions. Such reactions form networks of causal dependencies, in the sense that the outcomes of certain reactions form the conditions or inputs that trigger further reactions. Some of these networks turn out to be self-maintaining, in the sense that everything consumed (removed) by some reaction in the network is produced again by some other reaction, thus keeping the assembly of reactions running. Such a closed, self-maintaining network has been defined mathematically as a “chemical organization” (Dittrich and Fenizio 2007). It provides a simple, formal model for some of the most fundamental phenomena that we associate with life: self-organization, autopoiesis, cells, organisms, symbiotic assemblies and ecosystems.

In each of these cases, we see a higher level of organization emerging from the mutual adaptation of processes at the lower level. Thus, a stable agency, such as a cell, is in fact constituted by a self-perpetuating flux of reactions consuming and producing resources, such as molecules. This perspective, which sees stable wholes materialize out of processes, upends the one of the object ontology, which sees stable objects as primary and processes as secondary. Moreover, the reaction model defines a self-maintaining organization as fundamentally *autonomous*, initiating its own actions, rather than passively undergoing the effect of external forces. Thus, they are *teleonomic*, directed at the ultimate, intrinsic attractor of self-maintenance.

Of course, this does not mean that such organizations are independent of their environment. Even in the extreme case of the hermetically sealed ecosphere, thermodynamic constraints require an input of light and an output of heat. Thus, autonomous agencies must establish good relationships—i.e. dependable, synergetic

exchanges of resources—with their environment, and the other agencies that it contains.

While Chemical Organization Theory is well defined mathematically and computationally, its applications to biology are as yet limited. These include a number of preliminary, yet inspiring, models of the origin of life (Hordijk, Steel, and Dittrich 2018), metabolism (Centler et al. 2007; 2008), endosymbiosis (Veloz and Flores 2021), and ecology (Veloz 2019). For the wider ontological framework of relational agency introduced in this paper, the applications are even sketchier, and by necessity laid out in very broad strokes. Still, I have argued that this framework may provide a foundation for the highly desired “extended evolutionary synthesis” that would replace the too reductionist neo-Darwinist paradigm. This synthetic theory would broaden the assumed one-way effect of the environment on the selection of genes to a complex network of interactions between agencies at multiple levels, from molecules via genes, cells, and organisms, to colonies, ecosystems, and ultimately the planet.

While the complexity of such an integrated picture of the evolutionary process may seem daunting, I have tried to show with examples that the reaction network formalism allows one to start with very simple models, which can then be gradually refined and made more complex and realistic. Thanks to the power of computers, even complex models should remain manageable, in the sense that it would be possible to analyze which subnetworks of reactions form self-maintaining agencies, and which resources or species are likely to be eliminated (Centler et al. 2008). Developing such more realistic models will of course require much further research. Yet, I hope to have shown that, at the conceptual level at least, the ontology of relational agency already throws a new light on many old problems, such as self-organization, autopoiesis, symbiosis, teleonomy and multi-level selection, thus offering us a first glimpse of an eventual synthetic theory. I hope it will not be considered too presumptuous of me to see grandeur in this relational view of life, which not only builds further on the one of Darwin, but extends its horizons in both downwards and upwards directions, from particles to planets.

Acknowledgments

This research was funded by the John Templeton Foundation as part of the project “The Origins of Goal-Directedness” (grant ID61733), under its research program on “The Science of Purpose”. I thank my VUB colleagues collaborating on this project (Veloz et al. 2022), and in particular Shima Beigi, Tomas Veloz and Evo Busseniers, for many inspiring discussions on the ideas presented here.

References

- Abram, David. 2010. *Becoming Animal: An Earthly Cosmology*. Knopf Doubleday Publishing Group.
- . 2012. *The Spell of the Sensuous: Perception and Language in a More-Than-Human World*. Knopf Doubleday Publishing Group.
- Agafonov, Vladimir A., Vladimir V. Negrobov, and Abir U. Igamberdiev. 2021. “Symbiogenesis as a Driving Force of Evolution: The Legacy of Boris Kozopoliansky.” *Biosystems* 199 (January): 104302. <https://doi.org/10.1016/j.biosystems.2020.104302>.
- Ashby, W. R. 1962. “Principles of the Self-Organizing System.” In *Principles of Self-Organization*, edited by H. von Foerster and G. W. Zopf, 255–78. Pergamon Press. <http://csis.pace.edu/~marchese/CS396x/Computing/Ashby.pdf>.
- Astor-Aguilera, Miguel, and Graham Harvey. 2018. *Rethinking Relations and Animism: Personhood and Materiality*. Routledge.
- Baedke, Jan. 2013. “The Epigenetic Landscape in the Course of Time: Conrad Hal Waddington’s Methodological Impact on the Life Sciences.” *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 44 (4, Part B): 756–73. <https://doi.org/10.1016/j.shpsc.2013.06.001>.
- Barad, Karen. 2003. “Posthumanist Performativity: Toward an Understanding of How Matter Comes to Matter.” *Signs: Journal of Women in Culture and Society* 28 (3): 801–31.
- Bateson, G. 2000. *Steps to an Ecology of Mind*. University of Chicago Press.
- Bickhard, Mark H. 2011. “Some Consequences (and Enablings) of Process Metaphysics.” *Axiomathes* 21 (1): 3–32. <https://doi.org/10.1007/s10516-010-9130-z>.
- Bird-David, Nurit. 1999. “‘Animism’ Revisited: Personhood, Environment, and Relational Epistemology.” *Current Anthropology* 40 (S1): S67–91. <https://doi.org/10.1086/200061>.
- Bonabeau, E. 2002. “Agent-Based Modeling: Methods and Techniques for Simulating Human Systems.” *Proceedings of the National Academy of Sciences* 99 (Suppl 3): 7280–87. <https://doi.org/10.1073/pnas.082080899>.
- Braidotti, Rosi. 2013. *The Posthuman*. John Wiley & Sons.
- Camazine, S., J. L. Deneubourg, N. R. Franks, J. Sneyd, G. Theraula, and E. Bonabeau. 2003. *Self-Organization in Biological Systems*. Princeton Univ Pr.
- Centler, Florian, Pietro Speroni di Fenizio, Naoki Matsumaru, and Peter Dittrich. 2007. “Chemical Organizations in the Central Sugar Metabolism of *Escherichia Coli*.” In *Mathematical Modeling of Biological Systems, Volume I*, 105–19. Springer. http://link.springer.com/chapter/10.1007/978-0-8176-4558-8_10.
- Centler, Florian, C. Kaleta, P. S. di Fenizio, and P. Dittrich. 2008. “Computing Chemical Organizations in Biological Networks.” *Bioinformatics* 24 (14): 1611–18.
- Charlton, Bruce G. 2007. “Alienation, Recovered Animism and Altered States of Consciousness.” *Medical Hypotheses* 68 (4): 727–31. <https://doi.org/10.1016/j.mehy.2006.11.004>.
- Corning, Peter A. 2003. *Nature’s Magic: Synergy in Evolution and the Fate of Humankind*. Cambridge University Press Cambridge, UK.
- . 2022. “Teleonomy in Evolution: ‘The Ghost in the Machine.’” In *Teleonomy in Evolution*. Vienna Series in Theoretical Biology. MIT Press.
- Dawkins, Richard. 2006. *The Selfish Gene*. 3rd ed. Oxford University Press, USA.
- Dittrich, Peter, and Pietro Speroni di Fenizio. 2007. “Chemical Organisation Theory.” *Bulletin of Mathematical Biology* 69 (4): 1199–1231. <https://doi.org/10.1007/s11538-006-9130-8>.
- Fábregas-Tejeda, Alejandro, and Francisco Vergara-Silva. 2018a. “Hierarchy Theory of Evolution and the Extended Evolutionary Synthesis: Some Epistemic

- Bridges, Some Conceptual Rifts.” *Evolutionary Biology* 45 (2): 127–39. <https://doi.org/10.1007/s11692-017-9438-3>.
- . 2018b. “The Emerging Structure of the Extended Evolutionary Synthesis: Where Does Evo-Devo Fit In?” *Theory in Biosciences* 137 (2): 169–84. <https://doi.org/10.1007/s12064-018-0269-2>.
- Ferber, J., O. Gutknecht, and F. Michel. 2004. “From Agents to Organizations: An Organizational View of Multi-Agent Systems.” *Lecture Notes in Computer Science*, 214–30.
- Hartvigsen, G., A. Kinzig, and G. Peterson. 1998. “Complex Adaptive Systems: Use and Analysis of Complex Adaptive Systems in Ecosystem Science: Overview of Special Section.” *Ecosystems* 1 (5): 427–30.
- Harvey, Graham. 2014. *The Handbook of Contemporary Animism*. Routledge.
- . 2019. “Animism and Ecology: Participating in the World Community.” *The Ecological Citizen* 3 (1): 79–84.
- Heylighen, Francis. 1990. “Relational Closure: A Mathematical Concept for Distinction-Making and Complexity Analysis.” In *In: Cybernetics and Systems '90, R. Trappl (Ed.), (World Science, Singapore)*, p. 335-342.
- . 1999. “Advantages and Limitations of Formal Expression.” *Foundations of Science* 4 (1): 25–56. <https://doi.org/10.1023/A:1009686703349>.
- . 2001. “The Science of Self-Organization and Adaptivity.” In *The Encyclopedia of Life Support Systems*, 5:253–80. EOLSS Publishers Co Ltd. <http://pespmc1.vub.ac.be/Papers/EOLSS-Self-Organiz.pdf>.
- . 2011. “Self-Organization of Complex, Intelligent Systems: An Action Ontology for Transdisciplinary Integration.” *Integral Review*. <http://pespmc1.vub.ac.be/Papers/ECCO-paradigm.pdf>.
- . 2012. “A Tale of Challenge, Adventure and Mystery: Towards an Agent-Based Unification of Narrative and Scientific Models of Behavior.” ECCO Working Papers 2012–06. ECCO. Brussels, Belgium: Vrije Universiteit Brussel. <http://pcp.vub.ac.be/papers/TaleofAdventure.pdf>.
- . 2014. “Challenge Propagation: Towards a Theory of Distributed Intelligence and the Global Brain.” *Spanda Journal* V (2): 51–63.
- . 2016. “Stigmergy as a Universal Coordination Mechanism I: Definition and Components.” *Cognitive Systems Research*, Special Issue of Cognitive Systems Research – Human-Human Stigmergy, 38 (June): 4–13. <https://doi.org/10.1016/j.cogsys.2015.12.002>.
- . 2019. “Entanglement, Symmetry Breaking and Collapse: Correspondences between Quantum and Self-Organizing Dynamics.” *Foundations of Science*. <http://134.184.131.111/Papers/QM-Complexity.pdf>.
- . 2022. “The Meaning and Origin of Goal-Directedness: A Dynamical Systems Perspective.” *Biological Journal of the Linnean Society* submitted.
- Heylighen, Francis, and Shima Beigi. 2018. “Mind Outside Brain: A Radically Non-Dualist Foundation for Distributed Cognition.” In *Socially Extended Epistemology*, edited by Andy Clark, Spyridon Orestis Palermos, and Duncan Pritchard. Oxford University Press. <https://doi.org/10.1093/oso/9780198801764.003.0005>.
- Heylighen, Francis, Shima Beigi, and Tomas Veloz. 2015. “Chemical Organization Theory as a Modeling Framework for Self-Organization, Autopoiesis and Resilience.” ECCO Working Papers 2015–01. <http://pespmc1.vub.ac.be/Papers/COT-ApplicationSurvey-submit.pdf>.
- Heylighen, Francis, P. Cilliers, and C. Gershenson. 2007. “Complexity and Philosophy.” In *Complexity, Science and Society*, edited by J. Bogg and R. Geyer, 117–34. Oxford: Radcliffe Publishing. [https://doi.org/10.1002/\(SICI\)1099-0526\(199807/08\)3:6<12::AID-CPLX2>3.0.CO;2-0](https://doi.org/10.1002/(SICI)1099-0526(199807/08)3:6<12::AID-CPLX2>3.0.CO;2-0).
- Heylighen, Francis, and J. M. Dewaele. 2002. “Variation in the Contextuality of Language: An Empirical Measure.” *Foundations of Science* 7 (3): 293–340. <https://doi.org/10.1023/A:1019661126744>.

- Holland, John H. 2012. *Signals and Boundaries: Building Blocks for Complex Adaptive Systems*. Cambridge, Mass: The MIT Press.
- Hordijk, Wim, Mike Steel, and Peter Dittrich. 2018. "Autocatalytic Sets and Chemical Organizations: Modeling Self-Sustaining Reaction Networks at the Origin of Life." *New Journal of Physics* 20 (1): 015011.
- Ingold, Tim. 2006. "Rethinking the Animate, Re-Animating Thought." *Ethnos* 71 (1): 9–20. <https://doi.org/10.1080/00141840600603111>.
- Kauffman, S. A. 1993. *The Origins of Order: Self-Organization and Selection in Evolution*. Vol. 209. Oxford University Press New York.
- Kuhlmann, Meinard. 2000. "Processes as Objects of Quantum Field Theory." *Faye et Al*, 365–88.
- Laland, Kevin N., Blake Matthews, and Marcus W. Feldman. 2016. "An Introduction to Niche Construction Theory." *Evolutionary Ecology* 30 (2): 191–202. <https://doi.org/10.1007/s10682-016-9821-z>.
- Laland, Kevin N., Tobias Uller, Marcus W. Feldman, Kim Sterelny, Gerd B. Müller, Armin Moczek, Eva Jablonka, and John Odling-Smee. 2015. "The Extended Evolutionary Synthesis: Its Structure, Assumptions and Predictions." *Proceedings of the Royal Society B: Biological Sciences* 282 (1813): 20151019.
- Latour, Bruno. 1996. "On Actor-Network Theory: A Few Clarifications." *Soziale Welt*, 369–81.
- Lyman, R. Lee. 2004. "The Concept of Equifinality in Taphonomy." *Journal of Taphonomy* 2 (1): 15–26.
- Macal, C. M., and M. J. North. 2010. "Tutorial on Agent-Based Modelling and Simulation." *Journal of Simulation* 4 (3): 151–62.
- Marmodoro, Anna, and David Yates. 2016. *The Metaphysics of Relations*. Oxford University Press.
- Maynard Smith, John, and Eors Szathmáry. 1997. *The Major Transitions in Evolution*. Oxford University Press, USA.
- Mayr, Ernst. 1974. "Teleological and Teleonomic, a New Analysis." In *Methodological and Historical Essays in the Natural and Social Sciences*, edited by Robert S. Cohen and Marx W. Wartofsky, 91–117. Boston Studies in the Philosophy of Science. Dordrecht: Springer Netherlands. https://doi.org/10.1007/978-94-010-2128-9_6.
- Mayr, Ernst, and William B. Provine. 1980. *The Evolutionary Synthesis*. Vol. 231. Cambridge, MA: Harvard University Press.
- Miller, John H., and Scott E. Page. 2007. *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton, N.J: Princeton University Press.
- Milonni, Peter W. 2013. *The Quantum Vacuum: An Introduction to Quantum Electrodynamics*. Boston: Academic Press.
- Mingers, John. 1994. *Self-Producing Systems: Implications and Applications of Autopoiesis*. Springer Science & Business Media.
- Mossio, Matteo, and Leonardo Bich. 2017. "What Makes Biological Organisation Teleological?" *Synthese* 194 (4): 1089–1114. <https://doi.org/10.1007/s11229-014-0594-z>.
- Müller, Gerd B. 2007. "Evo–Devo: Extending the Evolutionary Synthesis." *Nature Reviews Genetics* 8 (12): 943–49. <https://doi.org/10.1038/nrg2219>.
- Okasha, Samir. 2005. "Multilevel Selection and the Major Transitions in Evolution." *Philosophy of Science* 72 (5): 1013–25. <https://doi.org/10.1086/508102>.
- Palsson, Bernhard. 2015. *Systems Biology*. Cambridge University Press.
- Peter, Stephan, and Peter Dittrich. 2011. "On the Relation between Organizations and Limit Sets in Chemical Reaction Systems." *Advances in Complex Systems* 14 (01): 77–96. <https://doi.org/10.1142/S0219525911002895>.
- Pigliucci, Massimo, and Gerd B. Müller. 2010. "Elements of an Extended Evolutionary Synthesis." In *Evolution—the Extended Synthesis*. The MIT Press. <https://doi.org/10.7551/mitpress/9780262513678.003.0001>.

- Prigogine, I. 1980. "From Being to Becoming. Time and Complexity in the Physical Sciences." *New York: Freeman, 1980.*
- Razeto-Barry, Pablo. 2012. "Autopoiesis 40 Years Later. A Review and a Reformulation." *Origins of Life and Evolution of Biospheres* 42 (6): 543–67. <https://doi.org/10.1007/s11084-012-9297-y>.
- Rubin, Sergio, Tomas Veloz, and Pedro Maldonado. 2021. "Beyond Planetary-Scale Feedback Self-Regulation: Gaia as an Autopoietic System." *Biosystems* 199 (January): 104314. <https://doi.org/10.1016/j.biosystems.2020.104314>.
- Schilthuizen, Menno, ed. 2008. "Life in Little Worlds." In *The Loom of Life: Unravelling Ecosystems*, 1–10. Berlin, Heidelberg: Springer. https://doi.org/10.1007/978-3-540-68058-1_1.
- Sourjik, Victor, and Ned S Wingreen. 2012. "Responding to Chemical Gradients: Bacterial Chemotaxis." *Current Opinion in Cell Biology*, Cell regulation, 24 (2): 262–68. <https://doi.org/10.1016/j.ceb.2011.11.008>.
- Sternberg, Shlomo. 2010. *Dynamical Systems*. Courier Corporation.
- Tumulka, Roderich. 2006. "On Spontaneous Wave Function Collapse and Quantum Field Theory." In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 462:1897–1908. The Royal Society.
- Turchin, Valentin. 1993. "The Cybernetic Ontology of Action." *Kybernetes* 22 (2): 10–30. <https://doi.org/10.1108/eb005960>.
- Veloz, Tomas. 2019. "The Complexity–Stability Debate, Chemical Organization Theory, and the Identification of Non-Classical Structures in Ecology." *Foundations of Science*, November, 1–15. <https://doi.org/10.1007/s10699-019-09639-y>.
- Veloz, Tomas, and Daniela Flores. 2021. "Toward Endosymbiosis Modeling Using Reaction Networks." *Soft Computing*, 1–10.
- Veloz, Tomas, Pedro Maldonado, Evo Busseniers, Alejandro Bassi, Shima Beigi, Marta Lenartowicz, and Francis Heylighen. 2022. "An Analytic Framework for Systems Resilience Based on Reaction Networks." *Complexity*. <https://researchportal.vub.be/en/publications/an-analytic-framework-for-systems-resilience-based-on-reaction-ne>.
- Veloz, Tomas, and Pablo Razeto-Barry. 2017. "Reaction Networks as a Language for Systemic Modeling: Fundamentals and Examples." *Systems* 5 (1): 11. <https://doi.org/10.3390/systems5010011>.
- Walsh, Denis M. 2018. "Objectcy and Agency: Towards a Methodological Vitalism." In *Everything Flows*. Oxford University Press.
- West, Stuart A., Roberta M. Fisher, Andy Gardner, and E. Toby Kiers. 2015. "Major Evolutionary Transitions in Individuality." *Proceedings of the National Academy of Sciences* 112 (33): 10112–19.
- Wilson, David Sloan. 1997. "Altruism and Organism: Disentangling the Themes of Multilevel Selection Theory." *The American Naturalist* 150 (S1): S122–34. <https://doi.org/10.1086/286053>.