Collective Consciousness Supported by the Web: healthy or toxic?

Beigi, Shima; Heylighen, Francis

Published in:
Computational Collective Intelligence

Publication date:
2021

Document Version:
Accepted author manuscript

Collective Consciousness Supported by the Web: healthy or toxic?

Shima Beigi[1][0000-0002-9552-9016] and Francis Heylighen[1][0000-0001-5823-5898]

[1] Vrije Universiteit Brussel
sbeigi@vub.be, fheyligh@vub.ac.be

Abstract. We define the noosphere as the conscious level of the web, where global conversations are being held about collective challenges. To understand its dynamics, we review three neuroscientific theories of consciousness: information integration, adaptive resonance, and global workspace. These suggest that conscious thoughts are characterized by a "resonant", self-maintaining pattern of circulating information. This pattern should be sufficiently stable to be examined and dependably stored, yet sufficiently plastic to adapt to new input. The self-organizing dynamics of ideas circulating on the web, however, may settle in an attractor that is too resistant to accommodate new information. This results in a closed, toxic form of collective consciousness exemplified by conspiracy theories. We review the global discussion of the COVID-19 pandemic to illustrate healthy and unhealthy forms of noospheric consciousness. We then argue for the need to promote the healthy form via the modelling of the dynamics of idea propagation and the dissemination of narratives promoting open conversation.

Keywords: noosphere, collective intelligence, consciousness, neuroscience, global brain, COVID-19, conspiracy theories

1 Introduction

Collective intelligence is the ability of a group of agents to solve problems together that the same agents working individually cannot solve [1]. Computational collective intelligence [2] further enhances that ability by harnessing the power of computers, networks and algorithms to support the sharing of information between the different agents in the group. An impressive application of such collective intelligence is Wikipedia, the global encyclopedia that covers nearly the whole of human knowledge. Such a vast network of structured knowledge could never have been assembled without the technology of the World-Wide Web, which allowed millions of users worldwide to collaborate easily and effectively, by adding, linking, and editing pieces of knowledge made available on a shared platform.

The basic mechanisms underlying such distributed cognition are the propagation of information between agents, the integration and processing of that information so as to extract coherent patterns, the filtering out of errors, and the consolidation of the resulting insights in a shared memory (such as Wikipedia), so that they remain available for later use. An important role here is played by the algorithms that control the routing, selection and processing of information. Inadequate algorithms may lead to valuable information being ignored, or poor-quality data being inappropriately relied on.

Yet, next to algorithms, the outcome of the process also depends on social interactions: who communicates what with whom? When people propagate information, e.g., by distributing an email, posting a message on social media, or making an edit in Wikipedia, they rely on their own judgment about what is appropriate to communicate. That judgment will depend on subjective factors, such as their feelings, personal beliefs, norms about public behavior, and relationships with others. These social and psychological variables affect the quality of the information being propagated, processed and stored, and therefore the collective intelligence of the overall human-computer network. To enhance the quality of the solutions produced, we need to better understand the dynamics of this immensely complex network of interacting agents.

The problem is that the number and diversity of all the human and computational variables involved is so large that it seems impossible to model this process in any detail. In previous research, we have therefore proposed to view this world-wide collective intelligence as a *global brain* [3, 4]. The idea is that we may better understand this network through its similarity with another extremely complex, self-organizing network, the human brain. In this analogy, individual agents (people, computers, websites, applications…) play the role of neurons in the brain. Messages being communicated between these agents are similar to electrical impulses or "activation" being transmitted between neurons via their connecting synapses. The links between agents—such as "friend" or "follower" links on social media, routing directions between computers or hyperlinks between web pages—play the role of these synapses. This "neural network" learns new knowledge when consolidated information is registered in a shared database or linking pattern.

In earlier work, we have investigated this model of the world-wide web as a global brain mostly at the level of algorithms that would create and exploit link structures to guide the propagation of information [1, 4]. This implements the brain mechanism of activation spreading in parallel along the strongest links. The idea is to support collective intelligence by connecting people or problems with the for them most appropriate solutions. Here, we wish to address the more complex issue of how new solutions, or more generally *new insights*, can be collectively generated through the synthesis of information from a variety of sources. This requires the more advanced brain mechanism of *conscious reflection*, in which different ideas and observations are critically examined, selectively accepted or rejected, improved and finally integrated into an overall interpretation that can guide further action.

We will illustrate this with the example of the ongoing global discussion about the causes of the COVID-19 pandemic and the best strategies to deal with them. To better understand how such a process takes place, we will introduce the concept of the *noosphere*, as the medium in which such discussions take place. We will briefly review contemporary theories of consciousness in the human brain, and suggest how they may be extended to the noosphere. This will provide us with a preliminary model for the self-organizing dynamics that characterizes the emergence of collective understanding. We will in particular distinguish healthy from unhealthy forms of collective consciousness, and suggest strategies to promote the former and prevent the latter.

2 Noospheric consciousness

The great majority of cognitive processes in the human brain are *subconscious*: they take place immediately and automatically in the background, leaving no occasion for the mind to examine precisely how or why the process came to a certain conclusion [5, 6]. For example, we know from neural network simulations that processing perceived sound into meaningful sentences is an extremely complex process, which involves thousands of variables, and which must pass through several stages or layers of abstraction. Yet, our brain does it effortlessly—although we have no idea how that happens, and we would not be able to prevent it from happening, even if the final interpretation turns out to be incorrect. The process also does not leave any trace in our memory: we may remember the final interpretation, but the intermediate stages in coming to that conclusion are not retrievable.

Another example of a subconscious process is the control of walking: we may consciously intend to go in a particular direction, but we are not aware of how our nervous system generates the precise muscular contractions that make our legs perform the right movements without losing our balance or swerving from the path. It is only when the situation deviates from our expectations—for example when we stumble on an obstacle we had not noticed—that we become conscious of the positioning of our body. That conscious focus is necessary to make sense of the novel challenge and to consider possible ways to deal with it. As long as the incoming information confirms our subconscious expectations, the brain assumes that it can trust those automatic processes, and that therefore there is no need for conscious reflection [5].

Computer-supported systems for distributed cognition, such as the World-Wide Web, follow this logic of automatic processes that are implicitly trusted to produce the right result. For example, when you send an email, you assume by default that it will reach the right address, and when you search for information on a well-known subject, such as tuberculosis, you assume that the page (e.g. in Wikipedia) you find will provide correct information. However, when a novel, unexpected challenge appears, such as the COVID-19 pandemic, there are no established routines or knowledge bases available that specify how to deal with it. Such a situation requires conscious reflection. But because the challenge is much too large and complex for any individual to solve, that reflection must be collective, distributed and global, i.e. it must take into account the widest possible range of people, observations and hypotheses. That requires an ongoing global exchange of ideas, opinions and data, a critical examination of those ideas, and a mechanism to integrate these ideas into an overall understanding and strategy.

Nearly a century ago, the visionary scientist and philosopher Teilhard de Chardin proposed the concept of the *noosphere* ("mind sphere"), which he conceived as a layer of thought that envelops the Earth, or as a collective brain constituted out of billions of individual brains communicating via a network of links [7, 8]. While this concept to some degree anticipated the World-Wide Web, it adds that this network would support a global reflection that would allow humanity to resolve its problems, and thus converge to a coherent understanding. Therefore, we can interpret the noosphere as the conscious level of the global brain, with the web as the infrastructure that supports the required propagation of information. To better understand the dynamics of the noosphere, we can draw inspiration from contemporary neuroscientific theories of consciousness in the brain.

3 Neuroscientific theories of consciousness

Attempts to define consciousness are copious. Yet, consciousness remains an open scientific challenge, subject to much debate and controversy. Historically, consciousness studies have been mostly centered on the study of subjective experiences, aka 'qualia'. The complexity of these experiences has stimulated the scientific study of the underlying neural mechanisms of consciousness, with Crick and Koch among the pioneers [9]. Given that there is no obvious equivalent of subjective experience in the noosphere, we will here focus on a more concrete aspect of consciousness that has been called "access consciousness" [10]. This refers to the ability to monitor, examine, register and, if necessary, redirect mental processes [5]. As we noted, the ability to access mental functioning is precisely what distinguishes a conscious process from a subconscious one, whose automatism prevents it from being examined, corrected or redirected.

The more recent neuroscientific theories of consciousness have been able to capture some of the complexity of consciousness without compromising on the need for scientific evidence or *a priori* alienating philosophical, mystical and spiritual views of consciousness. These theories include, but are not limited to, Information Integration, Global Workspace, and Adaptive Resonance. These theories, although different in their perspective, share the basic idea that the flow of information in the brain is key to understanding consciousness. In the following sections we expand on each theory.

3.1 Information Integration Theory (IIT)

IIT was developed by Giulio Tononi in order to address two of the most difficult challenges in investigating the generation of consciousness within/among brain regions. The first challenge is to specify the conditions that determine a brain area's role in generating conscious experiences. The second challenge is to determine the relative significance of certain regions of the brain, for example thalamocortical region in comparison to cerebellum [11]. Considering these two challenges, Tononi defines consciousness as "the capacity of a system to integrate information." This capacity is determined by the coherency of the underlying neural processes. Coherency measures the extent to which integration has reduced relative uncertainty (in the information theoretic sense [12]) between activities in different parts of the network.

IIT has become one of the key models for the underlying role of neurobiological factors in shaping subjective experiences. Moreover, because the mathematical measure Φ for information integration can in principle be applied to systems of any kind, it has initiated a discussion about the degree to which consciousness can be found in non-biological systems such as countries, networks of logic gates, or the Internet. On the other hand, the theory has been criticized for being too abstract and general, overlooking concrete neural conditions prerequisite for conscious experience, such as the binding together of sensory, motor and emotional activation patterns in the brain in the form of "semantic pointers" [13].

As research on consciousness progresses, it becomes increasingly evident that the subjectivity of conscious experience should no longer be seen as a purely mystical phenomenon [14]. It can rather be formulated as a synthetic phenomenon partly generated/influenced by neurophysiological factors and partly by individual differences in perception, learning styles, mental models and cultural backgrounds. This synthetic perspective has been framed as an opportunity to move towards a more systemic take on the study of consciousness [15].

3.2 Adaptive Resonance Theory

One of the criticisms of IIT is that it reduces consciousness to a static information measure, thus neglecting the complex dynamics of experience. In the Adaptive Resonance Theory (ART) developed by Stephen Grossberg, the emphasis is not initially on consciousness. Instead, ART positions itself as a cognitive neural theory that models the dynamics of the concrete processes within the brain, including categorization of information, learning, expectation, attention, synchronization, memory and search. These in turn support perception, cognition, and finally consciousness [16].

Attention regulates whether some new information is learned or ignored. Therefore it plays a key role in the ART approach. Learning takes place above a certain threshold determined by a vigilance parameter that controls the relative importance of bottom-up (incoming data) and top-down processes (expected patterns). One of the key issues addressed by ART is the *stability-plasticity dilemma*, which is well-known within both artificial and biological systems [17]. It refers to the need for a system to be sufficiently plastic so that it can integrate new knowledge, yet sufficiently stable to prevent the loss of previously learned knowledge. ART investigates the delicate balance between stability and plasticity so as to prevent the problem of catastrophic forgetting in artificial neural networks [18].

ART proposes that conscious states in the brain are characterized by activation circulating back and forth between different neural regions, and in particular between perceptual layers (incoming data) and higher-level concepts that explain the perceptions (interpretations, expectations). This self-reinforcing "resonance" between zones in the brain's neural network produces a pattern of activation that is sufficiently intense and persistent to create a clear focus of attention and to become registered into long-term memory, while remaining sufficiently fluid to adapt to novel input.

The stability-plasticity dilemma has direct applications to hybrid human-technological systems. The flexible coupling and decoupling of biological as well as artificial agents with their surrounding environment has become a key area of investigation in designing systems that that need to adapt to changing environments. However, the balance between being sufficiently adaptive to acquire novel information (plasticity), while dependably retaining older information (stability) remains elusive.

Achieving it is especially difficult when it comes to intervention strategies in times when regular patterns break down, like during disasters and systemic shocks [19]. The resulting confusion can result in pathological forms of adaptation. Too strong tendencies towards either stability or plasticity can lead to the formation of maladaptive patterns of understanding, such as the conspiracy theories that we will discuss further. In such a situation, constructive intervention becomes extremely difficult because the system tends to get locked into a polarized, black-and-white, view of reality, lacking the fluidity to deal with conflicting information. While the underlying mechanisms that primate brains use to deal with the Stability-Plasticity dilemma remain unclear, ART and the model of noospheric consciousness discussed in this paper both endeavor to find ways to offset the loss of valuable knowledge from the system, while increasing its resilience against perturbations.

3.3 Global Workspace

The last neuroscientific theory of consciousness we wish to review, the Global (Neuronal) Workspace, was proposed by Bernard Baars [20] and elaborated in particular by Stanislas Dehaene [6]. Its main idea is that neural activity becomes conscious when it is "broadcasted" across the *global workspace*, which is a central crossroads of neural connections in the brain. As such, this activity can be examined by different more specialized brain modules that otherwise have few connections. The global workspace functions like a public forum or shared medium where these modules, playing the role of the agents that constitute the "society of mind", can enter in a conversation so as to come to a consensus on what to do. This is analogous to the noosphere where the web functions as a medium for global conversation between all cognitive agents.

Broadcasting in the brain requires widespread and intense activation that "reverberates" or circulates across the workspace network. For a subconscious stimulus to become conscious, the initial activation must be non-linearly amplified until it crosses the threshold for "ignition" [6], when the circulating activity is strong enough to become self-sustaining and reach all parts of the workspace. That allows it to be maintained for a while in working memory, where it can now be consciously examined and processed by more specialized brain circuits. This theory confirms that conscious activity is characterized by widespread, coherent circulation, but adds that this allows it to be scrutinized and if necessary redirected, thus supporting the access perspective.

Table 1 summarizes IIT, ART and Global Workspace commonalities and offers possible avenues of application at the noospheric level.

Table 1. Implementing Neuroscientific theories of consciousness at Noospheric levels

Neuroscientific theories	Commonalities	Noospheric level implementation
Information Integration Theory	Coherence of process	Communication patterns
Adaptive Resonance Theory	Making sense of information	Remaining open/Learning
Global Workspace Theory	Collective interactions	Planetary conversation

4 Self-organization of dynamic patterns

The recurrent theme in the different neuroscientific theories we reviewed is that consciousness arises from activation circulating back and forth between different parts of the brain, thus making the activity across these regions coherent or integrated. Such "resonance" moreover stabilizes the pattern of activity long enough for it to be kept in a working memory or global workspace, where it can be examined by various more specialized "agents" that may selectively reinforce, add or subtract elements of the pattern. A sufficiently stable pattern will be transferred to long-term memory, where it remains stored indefinitely and can be recalled when needed. On the other hand, while circulating, the pattern should remain sufficiently plastic to quickly adapt to new inputs from perception, thought or memory.

The emergence of such a coherent, dynamic pattern can be modelled as a process of *self-organization*, i.e. the appearance of a global order out of local interactions [21]. For the brain, this process is typically modelled with neural network simulations, where the stable pattern can be seen as an attractor of the dynamics of activation propagating between neurons via the connecting synapses. However, such a model seems less appropriate for the noosphere, where what is circulating is not one-dimensional "activation", but a variety of ideas that belong to different categories or species.

A perhaps more accurate modelling language is *reaction networks*, a formalism inspired by chemical reactions [4, 22]. A reaction network consists of a set of "species": $S = \{a, b, c, \ldots\}$ and a set of reactions $R = \{r_1, r_2, r_3, \ldots\}$. A reaction transforms a combination of species (its input or "reactants"), which are necessary for the reaction to take place, into another combination (its output or "products") that are generated by the reaction. Here is an example of a simple reaction network:

$$r_1\colon a + b \rightarrow c;\quad r_2\colon c \rightarrow c + a + d;\quad r_3\colon d \rightarrow b \tag{1}$$

A reaction has a rate, which specifies how quickly it consumes reactants and generates products. This rate parameter plays a role similar to the weight of a connection in a neural network. As a reaction runs, the amount of its products increases, while the one of its reactants decreases. Dittrich [23] has shown that a sufficiently rich reaction network tends to settle into a particular subset of the participating species and reactions, which he called a "chemical organization". Such an organization is "closed", in the sense that no new species are added by reactions acting on the set. It is also "self-maintaining", in the sense that any species consumed by reactions will be produced by other reactions at least as much as it is consumed. That means that while the reactions consume and produce individual species, the organization as a whole remains stable, because any species removed is eventually reconstituted (as exemplified by network (1)). Such on-going self-(re)production is called *autopoiesis*.

An organization is an *attractor* of the dynamics imposed by the reactions: once the system gets into this regime, it will stay there. Still, an organization may be perturbed by the external addition or removal of species or reactions. This will generally result in a new organization, consisting of a different mixture of species and reactions. Yet, some organizations may exhibit strong resistance to change, in the sense that their reactions are able to neutralize such perturbations. They thus exhibit a form of persistent goal-directedness, returning again and again to their preferred state of autopoiesis [22].

An organization in this sense resembles an adaptive resonance in which self-reinforcing processes circulate between multiple components, while maintaining a stable pattern of activity. Resistant organizations are particularly stable. However, this means they may lack the plasticity necessary to adapt to new input. This may lead to a common pathology of the noosphere, in which ideas circulating through the socio-technological network form a self-reinforcing pattern that a priori rejects contradictory evidence. We will now illustrate such dynamics by examining how different interpretations of the COVID-19 pandemic have reverberated through the noosphere.

5 The impact of COVID-19 on noospheric consciousness

COVID-19, a disease caused by the SARS-CoV-2 virus, was first observed on 12 December 2019 in Wuhan, China. Shortly after, the city of Wuhan became the first epicenter of COVID-19 [24]. On 11 March 2020, the World Health Organization triggered the global health pandemic protocol, and issued a set of public health guidelines in the majority of affected areas. The resulting lockdown measures resulted in catastrophic losses of jobs, sources of income and last but not least increased systemic unemployment [25]. Stressors such as pandemics are often categorized as *systemic shocks*, since they drastically perturb the dynamics of the societal system. Such shocks trigger multilevel change and adaptive responses [26], but may offer the possibility of renewal, transformation and learning.

The noosphere as we defined it is the realm of collective reflection. Dysfunctional thought patterns and maladaptive strategies circulating in this space can result in a systemic loss of trust and the spread of dangerously false information. In this age of global interconnectivity, it is critical to be aware of the impact of thinking patterns and adaptive strategies on the propagation of information. In spite of the manifest economic and psychological impact of COVID-19 health measures on the social fabric, to our knowledge their influence on the noosphere has not been studied in any detail yet. Loss of loved ones, social isolation, stress and confusing information are just a few of the direct effects, with possibly a myriad of future impacts yet to be realized. These may radically shift the current dynamics of the noosphere towards either health or toxicity.

As the pandemic unfolded and measures taken triggered strong reactions from the population, different thought-sharing practices began to emerge. On the one hand, we saw a surge in scientific research, resulting in an unprecedented number of publications on COVID-19 being exchanged. This led to the formation of a new thought-sphere or noosphere discussing the various aspects of the virus, its propagation rate, its ability to evolve to more contagious variants, and last but not least the effectiveness of treatments [27]. On the other hand, the public began its own noospheric process of making sense of the situation. The great uncertainty along with the consequent fears for an unpredictable future, the need to belong to a community, the need for security, reduced trust in institutions, and global confusion led to a bottom-up surge of conspiracy thinking [28]. Let us examine these divergent noospheric dynamics in more detail.

At the onset of the pandemic, global society was worried and wanted to understand as quickly as possible the cause and remedies of this disease. Since the available information was highly scattered, fragmentary and ever-changing, the only way to make sense of it had to be distributed: harnessing the collective intelligence of millions of individuals, databases and processing algorithms connected by the Internet to achieve a coherent understanding. This triggered perhaps the largest noospheric process of global reflection yet. At present (summer 2021), a broad consensus seems to have emerged among scientists and international organizations about the strategy to follow: slowing down the spread of the virus through measures such as face masks, social distancing, and travel limitations, while preventing infection by inoculating as many people as possible with an array of vaccines that have proven their efficacy. These recommendations have been recorded in shared memory, e.g. on Wikipedia and governmental websites. Note that these recommendations have changed over the months as new observations were integrated. For example, the emerging insight that virus particles are spread primarily through the air emphasized the need for aeration, open-air meetings and face masks, while reducing the initial focus on disinfection and hand washing.

This global scientific discussion exemplifies a healthy balance between plasticity (incorporating new insights) and stability (repeatedly confirming and dependably recording the acquired insights). However, in parallel, a number of alternative narratives were circulating across the noosphere that were much less healthy. As an illustration, we will reconstruct one of these, a conspiracy theory centered on billionaire philanthropist and former Microsoft chairman Bill Gates [29].

The origins of this narrative lie in facts known before the pandemic: Bill Gates through his foundation promoted vaccination, warned a few years ago about the danger of a pandemic, directed a company that was collecting data, and funded pharmaceutical research. A more dubious idea that circulated is that vaccines are inherently dangerous. When the pandemic seemed to appear out of nowhere, while vaccines were being touted as the only long-term solution, many people felt a need to explain why this happened, while justifying their fear of vaccines. It was sufficient for some to suggest that Gates would have an interest in everybody getting vaccinated to see a coherent interpretation emerge: Bill Gates funded research to engineer the virus that caused the pandemic in order to convince otherwise reluctant people to get a vaccine that would collect data on them via an injected microchip. The different ideas supporting this theory mutually reinforce each other. The creation of the virus by Gates explains why he was able to predict the pandemic in advance, his involvement in Microsoft why he wants to collect data, his investment in pharmaceutical companies why he was able to get the virus and vaccine produced, the microchip why vaccines are dangerous [29].

In future research, we hope to model such mutual reinforcement with the help of a reaction network formalism. In this model, the “species” would denote circulating ideas, their concentration the number of people entertaining these ideas, and the reactions the dynamics by which people entertaining ideas on the input side of the reaction tend to convince themselves and others of the ideas on the output side—e.g. by posting messages affirming these ideas on social media. Below are some examples of conceivable reactions, where $x \rightarrow y$ can be read as “idea x explains, confirms, or makes plausible idea y”. (Note that in proper reaction network language, this should be written as

$x \rightarrow y + x$, because species x remains present after the reaction: it is not "consumed". But to simplify notation we have left out the repeated occurrence of x.)

desire to collect data → *need for microchip implantation*
need for microchip implantation → *desire to get people vaccinated*
desire to get people vaccinated → *intent to create virus* + *vaccination promotion*
intent to create virus → *prediction of pandemic*
intent to create virus + *pharmaceutical research* → *creation of virus* → *pandemic*
pandemic + *vaccination promotion* → *people getting vaccinated*
people getting vaccinated → *collection of data* → *desire to collect data*

Such a rudimentary model will further need to include the rate with which a reaction convinces people. If these rates are high enough, the system of ideas making up the conspiracy theory would become self-maintaining, i.e. being produced in sufficient amounts to keep the theory alive, in spite of forgetting, competition with rival systems of propagating ideas, and "perturbing" counterarguments that reduce belief in certain component ideas (i.e. that "consume" the species corresponding to these ideas). Using the software simulation developed in our research group [22], we may then try to determine the precise conditions under which this idea system would survive and grow.

A characteristic of conspiracy theories is that they are highly resistant to perturbations, because lack of confirming evidence, or even counterevidence, is interpreted merely as a confirmation that the conspirators (here, Bill Gates and confederates) are powerful enough to effectively disguise their activities. Another reason for the success of such a theory is that it is simple, unambiguous and concrete. For people who were not able to follow the much more complex scientific discussion, this narrative suddenly explains everything they had difficulty understanding, while providing them with a clear guideline: do not get vaccinated! Such a theory is not only incorrect, it is plain dangerous, because it allows the virus to continue infecting and killing people that refused the vaccination. Yet, conspiracy theories such as this continue circulating among millions of people on social media all across the world [28].

6 Promoting a healthy noospheric consciousness

"Resistant" ideas such as conspiracy theories shift the dynamics of noospheric consciousness towards maladaptive behavior and loss of resilience (i.e. ability to recover from a systemic shock). Addressing the impact of complex challenges such as the COVID-19 pandemic on global society requires a systemic level of collective intelligence—one that is aware of complex trade-offs such as the stability-plasticity dilemma, and in particular the danger of self-reinforcing toxic thought patterns.

More generally, we can define the health of the noosphere as its ability to make sense of the global situation, by integrating a wide range of observations, coming from diverse human and computational agents, and thus developing a coherent strategy to solve problems in a sustainable manner. This requires being sufficiently plastic (open to new information) to adapt quickly to changes, while providing guidelines that are sufficiently stable and clear to concretely depend on. We believe that a healthy noosphere requires the following critical abilities.

First, it should be capable of harnessing the bottom-up intelligence emerging from millions of individual agents while steering it towards collective understanding, knowledge, wisdom, mindfulness, and resilience. As such, it would reduce confusion, conflict and uncertainty, while contributing to the overall well-being of the world population. A variety of apps and web communities have appeared over the past years that provide people with emotional and social support while relieving their stress and anxiety by proposing advice and supporting practices, such as meditation, healthy lifestyles, and mutual help. Such creative noospheric initiatives were motivated by human compassion combined with technologies to heal noospheric sufferings and pains [30].

Second, a healthy noosphere should be capable of recognizing toxic thought patterns and preventing their spread before they contaminate collective consciousness. These thought patterns include conspiracy theories and related denials of plain evidence, e.g. of climate change. Most dangerous are closed worldviews that a priori reject non-conforming ideas while demonizing opponents and thus inciting potentially violent conflict, as exemplified by the ideology of the Islamic State terrorist group, and certain forms of populism, nationalism and fundamentalism. Developing strategies to prevent their spread requires further research on the psychological and social factors that promote such thinking [31]. It also requires an investigation into the algorithms that social media use to selectively propagate information to individuals depending on what they are most interested in. These algorithms tend to create "echo chambers" and "filter bubbles" in which people only see news that reinforces their existing opinions [32], thus closing them off from divergent ideas, while creating a strongly "resonant", autopoietic dynamics. Simulations of how resistant "organizations" [22, 23] of mutually reinforcing ideas self-organize and grow may help us to understand and regulate such dynamics.

Teilhard's conceived unification of thought in the Noosphere does not imply homogeneity nor reducing the uniqueness of individuals [8]; it rather means cultivating an unlimited dialogue and welcoming a diversity of perspectives, because different ideas can complement each other, thus accelerating collective learning and evolution. Our proposal for a healthy noosphere is one that cultivates planetary interconnectivity, an interlinked network of consciousness and information, and last but not least a convergence of individual actions towards the wellbeing of the common good.

This further requires the development and dissemination of worldviews and narratives that promote an open-minded, cooperative and sustainable society and that make people aware of the importance of interdependence and global conversation. A well-thought out example of such an approach is the noopolitik for the noosphere strategy proposed by Ronfeldt and Arquilla [33]. Another example is the recently founded *Human Energy Project* (HEP, https://humanenergy.io/) . Inspired by the work of Teilhard de Chardin and the concept of noosphere, the HEP is an initiative to develop a new planetary narrative called "the Third Story", which goes beyond traditional religious and reductionist worldviews to focus on the long-term evolution of life, mind and society. It in particular tries to reach young people, presenting them with an optimistic view of the future and a sense of meaning and direction that is consistent with science without being inconsistent with more spiritual beliefs. The project aims at fostering such awareness by clarifying what the noosphere is and how it functions, formulating strategies for its development towards an open, self-aware consciousness, while disseminating such insights towards the general public via publications, videos, the web and social media. The present paper can be seen as a contribution to these objectives.

Acknowledgments: this research was supported by the Kacyra Family Foundation as part of its Human Energy Project and by the John Templeton Foundation through its grant ID61733 "The Origins of Goal-Directedness".

References

1. Heylighen, F.: Collective Intelligence and its Implementation on the Web: algorithms to develop a collective mental map. Comput. Math. Organ. Theory. 5, 253–280 (1999). https://doi.org/10.1023/A:1009690407292.
2. Nguyen, N.T., Hwang, D., Szczerbicki, E.: Computational collective intelligence for enterprise information systems. Enterp. Inf. Syst. 13, 933–934 (2019). https://doi.org/10.1080/17517575.2019.1640394.
3. Heylighen, F.: From Human Computation to the Global Brain: the self-organization of distributed intelligence. In: Handbook of Human Computation. pp. 897–909. Springer (2013).
4. Heylighen, F.: The Offer Network Protocol: mathematical foundations and a roadmap for the development of a global brain. Eur. Phys. J. Spec. Top. 226, 283–312 (2017). https://doi.org/10.1140/epjst/e2016-60241-5.
5. Heylighen, F.: Mind, Brain and Body. An evolutionary perspective on the human condition. (2020).
6. Dehaene, S.: Consciousness and the Brain: Deciphering How the Brain Codes Our Thoughts. Penguin (2014).
7. Shoshitaishvili, B.: From Anthropocene to Noosphere: The Great Acceleration. Earths Future. 9, e2020EF001917 (2021). https://doi.org/10.1029/2020EF001917.
8. Teilhard de Chardin, P.: The Formation of the Noosphere. In: The Future of Man. p. 165. Image Books Doubleday, New York (1959).
9. Crick, F., Koch, C.: The Problem of Consciousness. Sci. Am. 267, 152–159 (1992).
10. Block, N.: On a confusion about a function of consciousness. Behav. Brain Sci. 18, 227–287 (1995).
11. Tononi, G.: An information integration theory of consciousness. BMC Neurosci. 5, 42 (2004). https://doi.org/10.1186/1471-2202-5-42.
12. Shannon, C.E.: A mathematical theory of communication. Bell Syst. Tech. J. 27, 379–423 (1948). https://doi.org/10.1002/j.1538-7305.1948.tb01338.x.
13. Thagard, P., Stewart, T.C.: Two theories of consciousness: Semantic pointer competition vs. information integration. Conscious. Cogn. 30, 73–90 (2014). https://doi.org/10.1016/j.concog.2014.07.001.
14. Lindhard, T.: Consciousness from the Outside-In and Inside-Out Perspective. J. Conscious. Explor. Res. 10, (2019).
15. Aru, J., Suzuki, M., Larkum, M.E.: Cellular Mechanisms of Conscious Processing. Trends Cogn. Sci. 24, 814–825 (2020). https://doi.org/10.1016/j.tics.2020.07.006.

16. Grossberg, S.: Adaptive Resonance Theory. In: Encyclopedia of Cognitive Science. American Cancer Society (2006). https://doi.org/10.1002/0470018860.s00067.
17. Grossberg, S.: The Attentive Brain. Am. Sci. 83, 438–449 (1995).
18. Mermillod, M., Bugaiska, A., Bonin, P.: The stability-plasticity dilemma: investigating the continuum from catastrophic forgetting to age-limited learning effects. Front. Psychol. 4, (2013). https://doi.org/10.3389/fpsyg.2013.00504.
19. Verbeke, P., Verguts, T.: Learning to synchronize: How biological agents can couple neural task modules for dealing with the stability-plasticity dilemma. PLOS Comput. Biol. 15, e1006604 (2019). https://doi.org/10.1371/journal.pcbi.1006604.
20. Baars, B.J.: Global workspace theory of consciousness: toward a cognitive neuroscience of human experience. In: Laureys, S. (ed.) Progress in Brain Research. pp. 45–53. Elsevier (2005). https://doi.org/10.1016/S0079-6123(05)50004-9.
21. Heylighen, F.: The science of self-organization and adaptivity. In: The Encyclopedia of Life Support Systems. pp. 253–280. EOLSS Publishers Co Ltd (2001).
22. Veloz, T., Maldonado, P., Busseniers, E., Bassi, A., Beigi, S., Lenartowicz, M., Heylighen, F.: An Analytic Framework for Systems Resilience Based on Reaction Networks. Complexity. (2021).
23. Dittrich, P., Fenizio, P.S. di: Chemical Organisation Theory. Bull. Math. Biol. 69, 1199–1231 (2007). https://doi.org/10.1007/s11538-006-9130-8.
24. Surveillances, V.: The epidemiological characteristics of an outbreak of 2019 novel coronavirus diseases (COVID-19)—China, 2020. China CDC Wkly. 2, 113–122 (2020).
25. Beigi, S.: How do the Covid-19 Prevention Measures Interact with Sustainable Development Goals? ArXiv201102290 Phys. Q-Fin. (2020). https://doi.org/10.20944/preprints202010.0279.v1.
26. Beigi, S.: THE COVID-19 PANDEMIC: A WAKEUP CALL TO ENHANCE SYSTEMIC RESILIENCE GLOBALLY. Wiley J. (2021).
27. Thanh Le, T., Andreadakis, Z., Kumar, A., Gómez Román, R., Tollefsen, S., Saville, M., Mayhew, S.: The COVID-19 vaccine development landscape. Nat. Rev. Drug Discov. 19, 305–306 (2020). https://doi.org/10.1038/d41573-020-00073-5.
28. Douglas, K.M.: COVID-19 conspiracy theories. Group Process. Intergroup Relat. 24, 270–275 (2021). https://doi.org/10.1177/1368430220982068.
29. Thomas, E., Zhang, A.: ID2020, Bill Gates and the Mark of the Beast: How COVID-19 Catalyses Existing Online Conspiracy Movements. Australian Strategic Policy Institute (2020).
30. Edwards, S.D.: AI in the noosphere: an alignment of scientific and wisdom traditions. Ai Soc. 1–3 (2020). https://doi.org/10.1007/s00146-020-00999-9.
31. Uscinski, J.E., Enders, A.M., Klofstad, C., Seelig, M., Funchion, J., Everett, C., Wuchty, S., Premaratne, K., Murthi, M.: Why do people believe COVID-19 conspiracy theories? Harv. Kennedy Sch. Misinformation Rev. 1, (2020). https://doi.org/10.37016/mr-2020-015.
32. Terren, L., Borge-Bravo, R.: Echo Chambers on Social Media: A Systematic Review of the Literature. Rev. Commun. Res. 9, 99–118 (2021).
33. Ronfeldt, D., Arquilla, J.: The Continuing Promise of the Noösphere and Noöpolitik: Twenty Years After. Social Science Research Network, Rochester, NY (2018). https://doi.org/10.2139/ssrn.3259425.