

UNIVERSAL ETHICS: ORGANIZED COMPLEXITY AS AN INTRINSIC VALUE

Clément Vidal
Center Leo Apostel &
Evolution Complexity and Cognition,
Vrije Universiteit Brussel,
contact@clemvidal.com

Jean-Paul Delahaye
Centre de Recherche en Informatique,
Signal et Automatique de Lille (CRISTAL)
Université de Lille
jean-paul.delahaye@univ-lille1.fr

ABSTRACT: How can we think about a universal ethics that could be adopted by any intelligent being, including the rising population of cyborgs, intelligent machines, intelligent algorithms, or even, potentially, extraterrestrial life? We generally give value to complex structures, to objects resulting from long periods of work, and to systems with many elements and with many links finely adjusted. These include living beings, books, works of art, and scientific theories. Intuitively, we want to keep, multiply, and share such structures, as well as prevent their destruction. Such objects have value not because more information (in bits) would simply mean more value. Instead, they have value because they require a long computational history – assuming that the numerous interactions governing their construction constitute a computation – and they display what we call *organized complexity*. To propose the foundations of a universal ethics based on the intrinsic value of organized complexity, we first discuss conceptions of complexity, and argue that Charles Bennett’s *logical depth* is a satisfactory notion of what we are looking for. We then put forward three fundamental imperatives: to *preserve*, *augment*, and *recursively promote* organized complexity. We show a broad range of applications with human, non-human, and non-living examples. Finally, we discuss some specific issues of our framework such as the distribution of complexity, the managing of copies and erasures, and how our universal ethics tackles classical ethical issues. In sum, we propose a clear, homogenous, and consistent ethical foundation that can integrate many universal ethics desiderata.

KEYWORDS: Philosophy of information, organized complexity, Kolmogorov complexity, logical depth, ethics, universal ethics, infoethics, thermoethics, environmental ethics, transhumanism

1. Introduction.....	1
2. Conceptions of complexity.....	4
3. Three imperatives.....	7
3.1 Preserve organized complexity.....	7
3.2 Augment organized complexity.....	8
3.3 Recursively promote organized complexity.....	10
4. Applications and examples.....	11
5. Discussion & Objections.....	12
5.1 Organized complexity and destruction.....	12
5.2 Organized complexity and traditional ethical issues.....	13
5.3 Anticipating some misunderstandings.....	14
6. Conclusion.....	16
7. References.....	16

1. Introduction

Humans have developed multiple kinds of ethical systems, in different parts of the world, based on religious or humanist values. In our global and post-modern era, moral relativism makes a strong case, but by definition, it is not attempting to find a foundation for ethics.

Most of the time, ethical systems are anthropocentric in the sense that they value human happiness above anything else. Furthermore, a well-founded normative theory should also be able to answer questions not only about human

values, but also other value-related questions, such as aesthetic ones, for example: “why does a symphony have value in itself?”

Globally, ethical systems from various cultures are often mutually incompatible, and when they are interpreted dogmatically, i.e. as the only right way to assign value and to act in the world, they generate conflicts and violence at social levels: communal, societal, national, and international.

In this chapter, we propose a new concept for the foundation of a universal ethics. By “universal ethics” we mean that it aims to be universally applicable by *any* valuating agent, be it a human, organization, robot, software agent, or extraterrestrial being. It also aims to be able to give value to *any physical object* in the universe. Furthermore, because of its mathematical definition, it could also apply to virtual, abstract, mathematical worlds (e.g. virtual reality, artificial life simulations, multiverse theories, mathematical proofs, etc.).

The “universal” declaration of human rights is not universal in this sense, as “universal” in the declaration refers simply to all human beings. It is unlikely that this declaration would make any sense to an extraterrestrial being in another galaxy, and it would be of no guidance to artificial agents in a virtual world. A purely anthropocentric ethics could also not say if it would be good or bad to annihilate a newly found exoplanet that was teeming with life.

Two kinds of universal ethical systems have been proposed. The first kind is based on *matter-energy* processes.

For example, *thermoethics*’ central principle is to make the most of free energy, and to avoid the production of unnecessary waste, disorder and entropy (it is also called “entropy ethics”, see Ostwald 1912; Freitas Jr 1979, sec. 25.1.3; Hammond 2005; Korbitz 2010; Vidal 2014, chap. 10).

Another example of a matter-energy path towards a universal ethics could be based on the concept of *emergy* (with an “m”). The concept comes from systems ecology, and entails a measurement of energetic content. It is defined as the value of a system, be it living, social, or technological, as measured by the solar energy that was used to make it (e.g. Odum 2007).

Although matter-energy universal ethical systems such as thermoethics remain underexplored, our chapter will not focus on the possibility of such systems. Instead, we will focus on a second kind of universal ethics, based on *information* and *computation*. We thus assume implicitly an ontology and metaphysics of information and computation (Delahaye and Vidal 2018). With the rise of the information society, and the importance and ubiquity of computers in our world, computation, information storage, and information exchanges are reshaping our selves and our societies. A philosophy based on information and computation, therefore, is becoming more and more relevant (for an introduction to the field, see Floridi 2003).

An information and communication approach to universal ethics can be based on *cybernetics*, as it is a general science of control and communication (for some steps in this direction, see e.g. Beer 1997; Chambers 2001; Vidal 2014, 285–86; Ashby 2017). Focusing on information science, Floridi (2008) developed the concept of *infoethics*, which bears similarities with thermoethics (Vidal 2014, 271). To deepen and broaden traditional ethical systems, Ward Bynum (2006) initiated the seeds of a universal ethics based on information that can apply to every physical entity in the universe. However, the concept of information used in Bynum’s approach is mostly *semantic*, and has been criticized for its

vagueness (e.g. Adriaans 2010). In this paper, we introduce and focus on a *syntactic* computational concept of valuable information, as the foundation of a universal ethics. This concept is *logical depth* (Bennett 1988). We think that this approach solves the issues that have been raised about infoethics: the basic notion is purely syntactic and mathematical, without the need to refer to cognitive agents such as humans. By contrast, the concept of semantic information that would be “well-formed, meaningful, and truthful data” is so delicate to define, and leads to so many difficulties, that it would risk obfuscating any ethical system based on it.

What are the desiderata of a universal ethics? A universal ethics should first be able to justify existing invariant values in humans, and provide a better understanding of why such commonalities exist. Empirical research has shown that there are broad invariants in human values: murder, theft, rape, lying, and destruction more generally are negative values in all societies, whereas health, wealth, friendship, honesty, safety, freedom, and equality are positive ones (e.g. Heylighen and Bernheim 2000).

A universal ethics should also be able to give value both to humans and their cultural products. These include works of art, scientific theories, books and libraries, and museums.

A universal ethics is also expected to support the widening of moral boundaries (see Fig. 1). We rarely see ethical concerns going beyond humanity. Fortunately, this is changing as consciousness is rising to care and value the non-living Earth’s climate, realizing that taking care of it is necessary at least for our long-term survival. Some religions such as Buddhism also do care about “all beings”, which may include all life on Earth and in the Universe.




Figure 1 - The widening of moral boundaries. With increasing awareness, humans care for increasingly wider systems. The self cares for the group, which extends from a family to a community, a nation, and eventually to the whole of humanity. We extend the last circles of compassion to the notion of complexity at large, not limited to living things.

Such widening of moral boundaries has taken several shapes, from the rise of animal rights (Bentham 1907, chap. XVII; Regan 1987), to taking care of all living things on Earth, for example with *biocentric ethics* (Agar 2001), *deep ecology* (Næss 2008), or *environmental ethics*. Indeed, the ecological worldview – according to which all living things and the Earth are connected – requires caring about non-human living processes that support the living realm.

In our digital era at the dawn of transhumanism, humans are transforming and re-shaping themselves through genetic engineering and technology. We are seeing the rise of the first augmented humans (or cyborgs), which raises entirely new moral issues that are hard to handle with traditional ethical systems. Whether we endorse transhumanism or not, we do need insights and frameworks to deal with current and future relations between humans and non-humans. A related field is *machine ethics*, which deals with moral aspects of machine-machine interactions (e.g. Anderson and Anderson 2011).

Finally, a genuinely universal ethics should lead consistently to an *extraterrestrial ethics*, i.e. an ethics that can meaningfully apply to potential extraterrestrial life (see Vakoch 2014). This field explores issues regarding our obligations and duties towards any new lifeform we may discover (see e.g. issues raised by planetary protection programs in Meltzer 2010). As a thought experiment, it is also challenging and mind broadening to think about the kinds of ethical behaviors and principles we can expect from potential extraterrestrial intelligence (see e.g. Ruse 1985), as a way to prepare for the impact of the discovery of extraterrestrial life (Vidal 2015).

In this paper, we first discuss various conceptions of complexity, and introduce the notion of *organized complexity*, based on the computational concept of *logical depth*. We then show how organized complexity can be put to use as an intrinsic value, leading to three core imperatives: that we should *preserve* and *augment* existing organized complexity, as well as *recursively promote* systems that increase organized complexity. We illustrate our framework with some applications and examples. Finally, we discuss various issues that arise from this original computational approach to universal ethics.

2. Conceptions of complexity

There is no doubt that multicellular organisms are more complex than unicellular ones. On large time scales, it is generally agreed that there is a growth in complexity of living beings that emerge from evolution (e.g. Coren 1998; Livio 2000; Mayfield 2013; Delahaye and Vidal 2018). Both biological and technological evolution produce increasingly complex objects, that is, displaying richer and richer structures.

A precise definition of such « complexity » is hard to achieve, and to formulate mathematically (for a review, see Bennett 1990). Many ideas have been attempted, but these have often been inadequate. For example, Eric Chaisson (2001) defines complexity as the rate of energy flowing through a system, normalized by its mass. This idea of energy rate density seems compelling at first, as it can apply to many epochs of cosmic evolution, from galaxy formation, planet formation, to living systems and our technological society.

However, it has the drawback that it is totally blind to the concepts of computation or information. There is no stable correspondence between energy and computation. As a matter of fact, a variant of Moore's law shows that since the 1960s the number of computations per unit of energy has been doubling every 1.5 years (Kooimey et al. 2011). A modern microprocessor is clearly more complex than a computer made 40 years ago, in the sense that for a given amount of energy, it can perform one million more operations. Structurally, a modern microprocessor also contains one million more transistors. Such aspects of complexity are not captured by a metric based on energy only.

Complexity can also rise in artificial life simulations, such as J. H. Conway's game of life (Berlekamp, Conway, and Guy 2001), in which energy plays no role. Another example is the complexity of a musical piece. It is not the piece's loudness, frequencies, or tempo that seem to determine its complexity, but rather its mathematical structure, without the need of any thermodynamic notion. Yet another example is a painting or a sculpture. According to Chaisson's metric, since no energy flows through them, they have zero complexity, and thus zero value if we were to take complexity as a guide for axiology.

Although useful to study ecosystem complexity (Ulgiati and Brown 2009), one can note that the concept of energy is not a good general proxy for complexity. For example, a lot of energy is necessary to produce an ingot of aluminum, but this doesn't mean that an ingot of aluminum is particularly complex. It thus seems valid to search for a concept of complexity that is *not* founded on energy and its circulation – to seek, that is, a non-thermodynamic concept of complexity.

In theoretical computer science, two main notions of complexity have been proposed: *Kolmogorov complexity* and Bennett's *logical depth* (Li and Vitányi 2008; Bennett 1988; Delahaye 2009). These notions concern finite numerical objects, that we can translate – without loss of generality – to finite strings s of '0' and '1'. The Kolmogorov complexity of s is by definition the *size* of the shortest program s^* that outputs the string s (for example, by printing it or by writing it in an output file). This is a useful notion in many respects, but it is not suitable as a measure of complexity if we consider complexity to be strongly structured. The notion does not cover the idea of complexification that scientists use when discussing biological or technological evolution. Indeed, strings that have the highest Kolmogorov complexity are random strings, and we intuitively know that randomness is the opposite of organized complexity.

Fortunately, based on Kolmogorov complexity, the notion of *logical depth* seems to be suitable as a definition of organized complexity, at least in a first approximation. The logical depth of s is by definition the *computation time* that it takes for the program s^* to produce s . In the case of an object s with low complexity (e.g. a repetitive string like 000000...00, or a random string) this computation time is minimal; whereas the more s can be regarded as a complex and structured object, the longer it takes for s^* to compute s .

For example, the string of the first million digits of π has a large logical depth, as well as the string of a musical piece translated into 0 and 1, or the string of prime numbers. The logical depth can either be seen as a measure of the quantity of structures in s , or as a quantity of computation present in s .

This definition of logical depth thus embeds two core values. First, because it is based on Kolmogorov complexity, it values the effort to find the shortest possible programs. This value mirrors the fundamental epistemic value

of searching for the simplest and shortest models in science. Given some data to explain, a simple and short theory is to be preferred over a longer one. This is known by philosophers of science as Occam's razor, but many computer scientists have formalized it using Kolmogorov complexity, showing the links between machine learning and compression (e.g. Blumer et al. 1987; Li, Tromp, and Vitányi 2003; Li and Vitányi 1992; 2008; Delahaye 1994).

The second value is that the longer it took for an object to appear (assuming that it cannot be obtained in a simpler way), the more value it has. In other words, the harder it would be to rebuild an object, the more value it has (Bennett 2014). This value is specific to the definition of logical depth, and its consequences will be explored in this chapter.

These computational concepts imply that we cannot naively consider the simple and the complex to be opposites. There are three notions at play: *simplicity*, *random complexity* (measuring the content in information via Kolmogorov complexity), and *organized complexity* (measuring the content in computation via logical depth). To augment organized complexity, one often needs to seek the simple: the simpler the computational mechanisms, the more efficient the production of computations will be. The simple is not opposed to organized complexity, it serves it and favors its growth.

Some aspects of normative evolutionary ethics valuing complexity growth and diversity are compatible with the ethics of organized complexity. For example, the survival instinct is clearly protecting the living complexity of the organism, while the reproductive instinct secures the preservation of genetic information through generations. Mutations and sexual reproduction, coupled with environmentally-induced selection, lead to a complexification of organisms on large time scales. The extent to which we should be inspired by evolutionary processes to build a normative ethics remains a huge debate that is well worth pursuing (see e.g. Maienschein and Ruse 1999; Quintelier, Speybroeck, and Braeckman 2011).

Evolution produces complex organisms by variation, inheritance and selection. Such elementary operations constitute what Dennett (1995) calls a « Darwinian algorithm », that we can assimilate to the creation of computational contents stored in living beings (see also Mayfield 2013). This algorithmic view can be naturally extended to cosmic evolution (Delahaye and Vidal 2018).

Complex organisms can only appear after less complex ones have already appeared. This corresponds to a *slow growth law*, that has been demonstrated for Bennett's logical depth (Bennett 1988). It shows that organized complexity cannot appear suddenly, and that it requires a long time of maturation. This can be illustrated with an example from the history of science, namely, the now-refuted idea of spontaneous generation (Strick 2000). From this computational perspective, spontaneous generation could have been refuted a priori, as it would be extremely improbable that sophisticated, complex living systems would appear after only a few days.

In sum, an ethical foundation based on this definition of complexity is conceivable. The good becomes what contributes to the preservation and augmentation of structural contents, in non-living entities, living beings, as well as humans and their cultural products, whether these contents are artistic, scientific, technological, economic, or political. Such an ethics suggests a respect and a protection not only of human and living beings, but also of all objects and structures that accompany them, or that further organizes them. Rich structures

have required extensive computational work, and this is why we must protect them.

A work of art (such as a painting, symphony, or novel) is an object with high computational content, as it is the result of extensive elaboration; its organization is rich and non-trivial. As an object with high logical depth, i.e. *computational content*, it has value and should be respected and protected, according to an ethics of organized complexity. The same holds for a science book, a microprocessor, the form of organization of our societies, the networks of interconnections of our cities and countries, etc.

It is remarkable that some cryptocurrencies such as Bitcoin have at their core a record - called the blockchain - that has high computational content, i.e. great logical depth. This makes them nearly unfalsifiable in practice, and constitutes a concrete example linking, in practice, the concept of logical depth with intrinsic value.

One may note that other ethics may be developed and supported with other existing or future conceptions of complexity, leading to somewhat different results, and this may be a worthwhile effort. However, it seems to us that the notion of logical depth lends itself naturally to the project of a non-thermodynamic universal ethics, as it offers a robust, adequate, and precise formal definition. We would like to show that the notion of logical depth encompasses and accommodates a wide range of existing values.

In what follows, we assume that logical depth captures satisfactorily the notions of richness in structure, high computational content, and organized complexity.

3. Three imperatives

We identify three imperatives of the ethics of organized complexity: to *preserve*, *augment*, and *recursively promote* what preserves and augments organized complexity. As these are fundamental imperatives, additional ethical consequences may be derived from them, but this is an exercise which we will not attempt to do systematically in this paper.

3.1 Preserve organized complexity

The first imperative is to *preserve existing organized complexity*. Indeed, organized complexity took time and effort to appear, so it makes sense to preserve and protect it. We have built-in biological survival instincts that lead us to preserve ourselves and our offspring. Humans also have a tendency to systematically collect and store organized complexity. For example, in recent times, more and more organizations have made efforts to preserve biodiversity, or to protect endangered species, and also to protect other kinds of things that have required great effort in order to exist: a painting, a monument, a patented idea. We are aware of the value of our cultural heritage, and that we should avoid the destruction of rare buildings, objects, or works of art. Historical and cultural preservation organizations exist worldwide, to preserve all kinds of complex human structures: consider for example UNESCO's world heritage sites, or many other national heritage protection programs.

In the digital world, this preservation imperative motivates us to implement effective strategies for backing up our data. Most of us have experienced data loss and know how costly this can be.

Even if the complexity collected has no immediate pragmatic value, it may have value in the future. This is true for biological evolution, where noncoding DNA sequences are conserved in the genome, and may in future generations be activated. This is also true in mathematics, where theories that once had no practical use are nowadays central tools for science and technology. Classical examples include the use of non-euclidean geometries for relativity theory, or the use of arithmetic for modern cryptography. A similar dynamic is likely to apply in the future, which strongly implies that we should carefully preserve and make accessible theoretical scientific knowledge.

If systematic collection and storage is not possible, we need to think about heuristics regarding what to collect, what to store, and what to make most accessible. In computer science, this is related to the recurrent problem of managing memory (space) and speed (time), which leads to space-time tradeoffs.

One may object that losing old *structural* products of complexity is not so grave, as long as we are able to *preserve their function*. For example, even if Galileo's original telescope design is not used anymore, we have much more powerful and reliable telescopes today, so the loss is limited. Another example is the computer. Nobody misses vacuum tube computers; what is important is that we have general-purpose computers that can accomplish the same operations as these earlier computers. The preservation of organized complexity, in other words, can be both structural and functional.

Also, as a first heuristic, one may suggest that we should adapt our preservation strategy relative to the number of copies. If there are 100 copies of a book, it is less grave to destroy one of them, than to destroy the last copy of a book.

3.2 Augment organized complexity

One can emphasize the preservation and conservation of the old, but the creation of the new, and the augmentation of existing organized complexity, are just as valuable and important. We saw that the slow growth law (Bennett 1988) implies that we can't just quickly create deeply complex objects out of nothing. We need time to build on previous efforts, which leads us to our second imperative: *augment organized complexity*.

This raises the question: how should we augment organized complexity? This is the issue of the *distribution of organized complexity*: should we try to augment total organized complexity, or average organized complexity? The issue is similar to classical discussion of the total versus average utility issue in utilitarianism (Sidgwick 1907). The issue is actually a particular instance of a more general problem of optimal allocation of resources, in defining social welfare functions (Chevaleyre et al. 2006).

To define the issue more precisely, let us imagine that we have a partition of the universe into well-defined components (humans, countries, celestial bodies, galaxies, etc.), that we denote as C_1, C_2, \dots, C_n . There are different goals that one may want to pursue, in order to preserve and augment organized complexity. For example:

Goal 1 : Augment organized complexity as a whole, as the union of: $C_1 + C_2 + \dots + C_n$.

Assuming that this organized complexity is measured adequately by logical depth (LD), the goal is thus to maximize $LD(C_1 + C_2 + \dots + C_n)$ that we note as:

maximize $(LD(C_1 + C_2 + \dots + C_n))$.

This may be called a *global* conception.

Goal 2 : Augment the sum of organized complexity inside the different components :

maximize $(LD(C_1) + LD(C_2) + \dots + LD(C_n))$.

This is not the same as Goal 1, because if C_1 and C_2 are identical, we have: $LD(C_1 + C_2) \approx LD(C_1) < LD(C_1) + LD(C_2)$. For example, if there are two identical books in one library they have almost the same value ($LD(C_1 + C_2) \approx LD(C_1)$). Striving for goal 2, the complexity of C_1 and C_2 will be counted in each component, and thus two times. The emphasis here is thus less on producing new complexity overall, but rather on distributing it evenly. In this context, we may call this goal 2 an *additive* conception.

Goal 3 : Maximize the least structurally complex components:

maximize $(\text{minimum}(LD(C_1), LD(C_2), \dots, LD(C_n)))$.

This is an *egalitarian* conception, aimed at making the poorest the least possible.

Goal 4 : Maximize the organized complexity of the best components :

maximize $(\text{maximum}(LD(C_1), LD(C_2), \dots, LD(C_n)))$.

In this case, it is an *elitist* conception : it doesn't matter if some components have low organized complexity, what matters is increasing the organized complexity of the already-highest complexity components.

Goal 5 : Maximize the product of organized complexity :

maximize $(LD(C_1) \times LD(C_2) \times \dots \times LD(C_n))$.

This is a compromise between goal 3 and goal 4 that avoids over-penalizing certain components. This conception is used in the optimal allocation of resources and is called *Nashian* (Ramezani and Endriss 2010).

We will not discuss how one might settle these different goals and viewpoints, as each of them has arguments in its favor. What we do want to argue is that the richness of a component can be measured by its organized complexity, and that this is a universal and coherent way to approach this issue.

Let us note that as soon as a metric to measure the value of components in the world is given, in any valuation system, this allocation of resources issue will arise, and we will need to choose between a global, additive, egalitarian, elitist, or Nashian allocation system. Even if the measure is not clear-cut or possible in practice, the different goals underlie different philosophical, political, and ethical choices. To decide between the different goals requires the

development of an applied ethics. Therefore this issue is not a weakness of our proposal in particular, as any foundational principle for ethics (e.g. maximizing human happiness) needs to decide between such goals when put into practice.

It is also worth noting that informational resources are non-rival, and can be shared with negligible costs compared to rival, matter-energy resources. So different strategies and treatments might be necessary for the distribution of rival and non-rival resources. For example, it is natural and easy to be egalitarian and to share knowledge with all humans via the internet, whereas doing the same with oil is much more problematic.

3.3 Recursively promote organized complexity

The third imperative is to *recursively promote what preserves and augments organized complexity*. In other words, it is to create, value and assist systems and strategies that can preserve and augment organized complexity, to the n^{th} order.

Let us illustrate this imperative with two examples: the mathematician and the musician. At the recursive level 0, the mathematician finds a new theorem, and the composer composes a new piece of music. They can make efforts to diffuse their works, for example, by writing a book or recording a CD. This is a level 1 effort, as it limits the risk that the organized complexity created could simply disappear. A level 2 effort, for example the funding or founding of a multimedia library, will allow the preservation and diffusion of organized complexity when the book or CD is released. A level 3 effort would include, for example, participating and helping a government whose goal was to collect taxes in order to fund the building of libraries.

One can note that a library security system does not have a strong intrinsic complexity, but is still a valuable aspect of preserving existing organized complexity, recursively. Jacques Monod (1972, 180) defended such a higher-order way to preserve knowledge, with his *ethic of knowledge* that “prescribes institutions dedicated to the defense, the extension, the enrichment of the transcendent kingdom of ideas, of knowledge, and of creation”.

Generally, copies also help realize this recursive imperative. For example, libraries that store copies of books help to further build complexity, as they give access to existing deeply complex work, on which further complexity can be built. In our digital era, it becomes obvious that all kinds of open source and open access initiatives are cheap and highly beneficial, and should therefore be promoted (e.g. Heylighen 2007; Steele 2012).

There are indirect ways to promote organized complexity, namely to ensure that its supporting systems are effective. For example, we need to care about Earth’s climate for the preservation and augmentation of biological complexity. Earth’s climate is non-human and non-living, but it should still be taken care of. Another example is the requirement of energy to build organized complexity. Energy has value in the sense that it could potentially be used to build organized complexity, which ties in with the values of thermoethics. Unfortunately, these recursive and indirect ways to promote organized complexity make the assessment of value more complicated in practice.

4. Applications and examples

If one adopts organized complexity as a universal value, then it becomes possible to naturally recover a large number of values that are already accepted by many ethical theories. Let's consider a few examples.

Every human being is a complex construct resulting from one's genes, learning, and experiences. Each individual human construction is unique and has a value of logical depth which, even if we do not know how to measure it precisely, is clearly very high. To kill a human being is, from the point of view of the ethics of organized complexity, a bad action. Similarly, anything that degrades, disturbs, or hurts a human being, by making it less effective and simply by damaging its structural richness, must be recognized as bad from this point of view.

The ethics of organized complexity commands us, in the same way, to respect and protect animals and, in a general way, all living things. Interestingly enough, this ethics gives special importance to endangered species, as their members carry an organized complexity that would be impossible to recover should the species disappear. The difference of care that should be devoted to a member of a species represented by millions of other members, as opposed to just a few, is naturally taken into account by the ethics of organized complexity. This idea is far from new. For example in the biblical myth of Noah's ark, God commands Noah to protect at least one sexual couple of each existing species. In this way, Noah not only *preserves* existing biological complexity, but thanks to the instruction to protect both sexes, also enables their reproduction and thus the *recursive promotion* of complexity.

In the case of human beings, a strict application of the ethics of organized complexity would at first sight lead to assigning more value to a genius than to the average person. However, we generally assume that all human lives have equal value. To suppose the opposite would seem to create serious social and political difficulties, and the solution that seems most compatible with democracy, and hence global social efficiency would be to regard all humans as having equal rights.

Another example of the immediate application of the ethics of organized complexity concerns works of art. These have obvious structural content, and the recognized talent of the artists who produce them is linked to their ability to elaborate (to calculate) complex, original structures in novel ways. Such creativity is a form of logical depth. Even if this computational content is not apparent in the work itself, when regarded as an isolated object – as with a work of minimalist art, for example – this content may still be present in the new relations which it establishes between the world and the work of art.

Still, generally speaking, we recognize that a work of art has value in proportion to its internal structural richness and its novelty. When it offers a novel perspective on the world, we assign it value in proportion to the fineness and subtlety of what it implies about the world. This value represents a new form of structural wealth, established by the work of art. In most cases, what makes us recognize value in a work of art can be interpreted as inherent structural richness, and therefore the idea that works of art must be preserved and protected is a direct application of the principles of the universal ethics of organized complexity.

Music is a particularly striking example of the purely structural content we perceive, and humans value many different kinds of music styles. While the development of our musical preferences will make us appreciate differently European or Indian music, for example, with effort we can learn to appreciate even music to which we have not been accustomed. In the end, what we love about music, and what makes a musical work worthwhile for us, is its richness in structure.

Science, too, can fall within the value scheme proposed by the ethics of organized complexity. The most important scientific theories – that would be most morally condemnable to forbid, or to not diffuse – are those which required greater effort (experimental, conceptual, mathematical, etc.). Of course, by “effort”, we mean a well-formed and well-informed effort. Those grand, deeply complex works have a great content in computation or, equivalently, in structure. Here our proposition of universal ethics uses its single, homogeneous concept to accommodate a fundamental value of scientific practice.

5. Discussion & Objections

5.1 Organized complexity and destruction

A delicate question that arises, if we adopt organized complexity as a value, is the erasure of data or, equivalently, the destruction of structures. Destruction seems to directly oppose the promotion of organized complexity. Yet, at least three factors determine the “right” decision to make in concrete cases.

First, the existence or non-existence of other copies of the data or structure in question. If there are many copies, erasing a few redundant copies will not affect the global organized complexity. One may thus argue that their destruction is not (so) “bad”.

However, as we saw, there is a second factor one may want to include, if one adopts a principle of egalitarianism for managing organized complexity. In this case, destroying complexity at one point could have the effect of lowering the organized complexity of the component considered, and would thus be “bad”.

A third difficulty arises when taking into account our third imperative of recursively promoting organized complexity. This recursive promotion often requires us to keep, at least temporarily, some data or structures. The problem is similar to that of information management in a computer system: efficiency sometimes requires that the same data be copied several times, in order to have optimal access to it, and thus to compute the desired results more quickly. Here, even if it is not immediately useful to keep the intermediate data or structures in question, because they do not contribute directly to the final desired result, it is possible that keeping copies of such data guarantees a better creative potential and therefore that the “good” choice is to destroy nothing.

Another aspect of this algorithmic issue is that some data are easy to reconstruct, and thus keeping them clutters the space or memory of the system. Good management then requires the destruction of such data. For example, modern algorithms for testing the primality of a number, or for the fast generation of prime numbers, make it unnecessary to build and store large tables of prime numbers. However, it made sense to do so in a world without computers, and indeed, in the nineteenth century, the building of such tables

gave rise to important works and publications. In this case, the technological progress of computing has tipped the right decision from "keep" to "delete".

In sum, even by adopting a mathematically well-defined point of view for assessing value, the precise determination of actions recommended by an ethics of organized complexity does not lead in practice to simple solutions. On the contrary, it preserves all of the difficulties that are necessary for conceiving efficient algorithms.

5.2 Organized complexity and traditional ethical issues

Let us outline a few connections between our proposal and some traditional ethical issues.

It may be possible – and desirable – to consider the rise of organized complexity as the utility function in a utilitarian framework (instead of, for example, maximizing human happiness). Of course, it remains very hard to foresee what actions will preserve and augment complexity, especially if we consider our third principle of recursively promoting organized complexity.

However, this issue is not specific to our approach, as any consequentialist must have good models of the world, to anticipate the future as well as possible. This leads to the cognitive value of modeling, i.e., the anticipation of our actions and the world. On a short time scale, we can all agree that having the reflex to avoid a falling rock, or the skills to climb a tree to collect a delicious fruit, would be “good” and adaptive. We need similar skills at higher spatial and temporal levels, to be able to anticipate and model the future of groups, societies, Earth, the universe, and complexity at large.

Instead of taking organized complexity as an intrinsic *normative value*, one could also use it as a *descriptive value*, in the sense that it may help to explain why an agent values certain behaviors or objects. One may thus try to explain an existing moral valuation in terms of its preservation and augmentation of organized complexity.

We can recover the value of striving for happiness, but not put it at the foundation of our ethics. For example, pleasures and pains evolved to ensure survival and reproduction, i.e. to preserve the organized complexity of our selves. Of course, we acknowledge that there are many reasons to promote happiness: for example, stress diminishes fitness for reproduction (see e.g. Moberg 1985), whereas happy people have positive emotions that allow them to broaden and build their selves, relationships, and environment (e.g. Fredrickson 2004).

Unfortunately the problem of assessing logical depth belongs to the class of non-computable problems. This means that we would rarely be able to prove that a fixed number is the correct logical depth value of a specific s , and that instead we must be content with algorithms that approximate $LD(s)$. However, this non-computability is not a fundamental obstacle to its use, as with Kolmogorov complexity, which is also non-computable is still widely used for concrete applications (e.g. Varré, Delahaye, and Rivals 1999; Belabbes and Richard 2008). So we can reasonably hope that similar tools can be used to approximate logical depth (for early attempts, see e.g. Zenil, Delahaye, and Gaucherel 2012; Gauvrit et al. 2017).

Another issue is the measurement of the depth of an isolated object. To what do we compare its organized complexity? For example, from what point in the past do we consider the history of our object? To assess the computational

content of a human being, do we say that it starts with its birth, or its parent's birth, or should we go back to the origin of life, or even to the origin of the universe that allowed its atoms to exist? This is the *temporal boundary issue* that is not specific to our approach, as it also appears in holistic approaches, such as the idea of *emergy* (Hau and Bakshi 2004, 221a): should we take into account all the solar energy that has been used since the birth of the solar system to assess a content in solar energy of an object? If not, from when do we start?

In practice, we can't yet compute the value of any choice to solve systematically ethical issues. Instead a multiple-level ethics requires us to solve the question: *What action preserves and augments complexity at all levels?* For example, deciding on whether to abort a fetus or not may require multiple problem-solving considerations at multiple levels, including biological, psychological, familial, societal, or religious aspects. Often, philosophers emphasize a central moral conflict: that between the interests of the individual and the collective. But this conflict need not to be limited to just two levels; as we grow our circles of compassion, many more levels must be included (see Fig. 1). Again, this delicate issue regarding multiple levels is not specific to the ethics of organized complexity, but is a general problem in any complex issue, ethical or not, that involves multiple levels, aspects, and stakeholders.

5.3 Anticipating some misunderstandings

5.3.1 Intentionality

To determine if an action is good or bad is a question of intentionality, why did you not address this issue?

The issue of intentionality is separate. For example, if X kills Y, in all ethical views that value human lives, this is bad, whether the act of X was intentional or not. The question of determining the responsibility of X, how X should be punished, is a different one than knowing if killing Y was good or bad. The problem of moral assessment is different from the problem of determining responsibility. The latter has to do with juridical and penal domains of knowledge. To put X in jail because he was not careful could be justified, but if Y committed suicide by jumping in front of the train that X was driving, this would make no sense.

5.3.2 Politics

How could we compare a dictatorial and a democratic regime, if both have about the same complexity?

To evaluate them, we would need to evoke the three imperatives *together*. We must analyze the capacity of these societies to produce more complexity in the future (second and third imperatives). If it was established that totalitarian societies were better at preserving organized complexity, and steering the creation of new complexity at all levels, then they should be preferred. However, it seems that the opposite is true, as democracies seem much more favorable to diversity, the flourishing of people, art, and science.

5.3.3 Applicability

Isn't this universal ethics too abstract and without any possible application?

Applications of ethical principles are delicate steps to take, as we saw for example with the problem of the distribution of organized complexity. Any ethical approach that wants to become precise and applied must face such issues.

For example, if one considers that *human happiness* is the ultimate good, this raises many questions about how to apply it in practice. Do we want to maximize the happiness of the most unhappy? Or the average or maybe the median happiness? Do we want to maximize the sum of happiness on Earth? Should we aim at making many humans averagely happy? Or should we aim at less happy humans, but more so on average? This is without mentioning the problem of measuring happiness.

In any ethical system, there are countless difficulties to bridging the gap between founding principles and determining real life actions. We have proposed a unified and mathematically based foundation that may ease this transition, although it remains a difficult problem.

5.3.4 Tastes and complexity

Isn't your proposal inaccurate, as some of our tastes show that value is not linked to complexity?

We did not argue that our tastes are always linked to what is the most complex; this is clearly wrong. Following one's values often demands effort. For example, one can prefer an airport novel to the hard work of a Nobel prize-winning novel, but if one had to destroy the last remaining copy of one or the other, one could still prefer to keep the Nobel prize-winning novel. One could also feel that to indulge in the consumption of the easiest-to-read novels would not always be the ethically best choice.

5.3.5 Ethical value and market value

Doesn't the market value of art –and minimalist art in particular– show that complexity has nothing to do with value?

Market value is generally not linked to the value of organized complexity (even if it may happen). The market value of a work of art and of goods in general depends on many parameters, such as rarity, tastes of buyers, etc. This divergence between market value and ethical value exists in all value systems: a lethal weapon will rarely have a great ethical value, and yet could be expensive.

5.3.6 Murderer and society

Shouldn't even a murderer, because of his biological structural complexity, be considered as good and worth respecting?

Of course not, because this first analysis is too quick and insufficient: it is restricted to only one level (the individual's complexity), and uses only the first imperative: "preserve organize complexity". Again, the three principles work *together*, and sometimes to preserve might be less important than to augment or to promote recursively organized complexity.

A murderer has killed another human being and therefore destroyed organized complexity. The murderer's potential future actions also put in danger the normal functioning of society and other human beings. Leaving the murderer free would create further societal problems, stresses and chaos, while executing the murderer – as it would have been done in the past – is not the best solution either, at least because it destroys the structural complexity of the murderer as a human being. So the problem is multi-level, and becomes one of preserving the complexity of *both* the murderer and society. It is generally at both of these levels that modern societies try to find solutions.

6. Conclusion

We have explored a possible foundation for ethics, by showing that organized complexity can be treated as an intrinsic and universal value. Such a non-anthropocentric, universal ethical foundation is much needed in our digital era. We have argued that organized complexity measures an intrinsic value: the history of non-trivial steps that have occurred to produce an object.

We put forward three imperatives of the ethics of organized complexity: to *preserve*, *augment*, and *recursively promote* what preserves and augments organized complexity. There are still many difficulties that lie ahead to apply those imperatives in practice. However, we saw that many such difficulties also exist in other ethical theories, and thus do not constitute a weakness of our approach in particular.

The potential of such a universal ethics is great, as it could be used to develop transhumanist ethics, machine ethics, or extraterrestrial ethics. It validates many existing notions of good and bad, such as the value of endangered species or works of art. An original non-anthropocentric conclusion is that deeply complex, inanimate objects also have value. In sum, we expect that the development of the ethics of organized complexity will shed light on past, present, and future ethical issues.

7. References

- Adriaans, Pieter. 2010. "A Critical Analysis of Floridi's Theory of Semantic Information." *Knowledge, Technology & Policy* 23 (1–2): 41–56. doi:10.1007/s12130-010-9097-5.
- Agar, Nicholas. 2001. *Life's Intrinsic Value; Science, Ethics, and Nature*. New York: Columbia University Press.
- Anderson, Michael, and Susan Leigh Anderson. 2011. *Machine Ethics*. Cambridge University Press.
- Ashby, Mick. 2017. "Ethical Regulators and Super-Ethical Systems." In *61st Conference of the International Society for the Systems Sciences*. Vienna, Austria. <http://ashby.de/Ethical%20Regulators.pdf>.
- Beer, Stafford. 1997. "The Culpablist Error: A Calculus of Ethics for a Systemic World." *Systems Practice* 10 (4): 365–80. doi:10.1007/BF02557886.
- Belabbes, Sihem, and Gilles Richard. 2008. "Spam Filtering without Text Analysis." In *Global E-Security*, edited by Hamid Jahankhani, Kenneth Revett, and Dominic Palmer-Brown, 144–52. Communications in Computer and Information Science 12. Springer Berlin Heidelberg. doi:10.1007/978-3-540-69403-8_18.
- Bennett, C. H. 2014. *Evidence, Computation, and Ethics*. Simons Symposium on Evidence in the Natural Sciences.
- Bennett, C.H. 1988. "Logical Depth and Physical Complexity." In *The Universal Turing Machine: A Half-Century Survey*, edited by R. Herken, 227–257. Oxford University Press. <http://researcher.watson.ibm.com/researcher/files/us-bennetc/UTMX.pdf>.
- . 1990. "How to Define Complexity in Physics and Why." In *Complexity, Entropy, and the Physics of Information*, edited by Wojciech H. Zurek, 137–48. Santa Fe Institute Studies in the Sciences of Complexity, v. 8. Redwood City, Calif: Addison-Wesley Pub. Co.

- Bentham, J. 1907. *An Introduction to the Principles of Morals and Legislation*. Oxford: Clarendon Press.
<http://www.econlib.org/library/Bentham/bnthPML0.html>.
- Berlekamp, Elwyn R., John Horton Conway, and Richard K. Guy. 2001. *Winning Ways for Your Mathematical Plays*. 2nd ed. Natick, Mass: A.K. Peters.
- Blumer, Anselm, Andrzej Ehrenfeucht, David Haussler, and Manfred K. Warmuth. 1987. "Occam's Razor." *Information Processing Letters* 24 (6): 377–380.
<http://www.cse.buffalo.edu/~hungngo/classes/2008/694/papers/occam.pdf>.
- Bynum, Terrell Ward. 2006. "Flourishing Ethics." *Ethics and Information Technology* 8 (4): 157–73.
 doi:10.1007/s10676-006-9107-1.
- Chaisson, E. J. 2001. *Cosmic Evolution: The Rise of Complexity in Nature*. Harvard University Press.
- Chambers, Jean. 2001. "A Cybernetic Theory of Morality and Moral Autonomy." *Science and Engineering Ethics* 7 (2): 177–92. doi:10.1007/s11948-001-0039-1.
- Chevaleyre, Yann, Paul E. Dunne, Ulle Endriss, Jérôme Lang, Michel Lemaitre, Nicolas Maudet, Julian Padget, Steve Phelps, Juan A. Rodriguez-Aguilar, and Paulo Sousa. 2006. "Issues in Multiagent Resource Allocation." *Informatica* 30 (1).
<http://wen.ijs.si/ojs-2.4.3/index.php/informatica/article/viewFile/70/62>.
- Coren, Richard L. 1998. *The Evolutionary Trajectory: The Growth of Information in the History and Future of Earth*. The World Futures General Evolution Studies, v. 11. Amsterdam, Netherlands: Gordon and Breach Publishers.
- Delahaye, Jean-Paul. 1994. *Information, complexité et hasard*. Paris: Hermes Science Publications.
 ———. 2009. *Complexité Aléatoire et Complexité Organisée*. Sciences En Questions. Versailles: Quae.
- Delahaye, Jean-Paul, and Clément Vidal. 2018. "Organized Complexity: Is Big History a Big Computation?" *American Philosophical Association Newsletter on Philosophy and Computers* 17 (2): 49–54. <http://arxiv.org/abs/1609.07111>.
- Dennett, Daniel Clement. 1995. *Darwin's Dangerous Idea: Evolution and the Meanings of Life*. New York: Simon & Schuster.
- Floridi, L., ed. 2003. *The Blackwell Guide to the Philosophy of Computing and Information*. Blackwell Publishing.
- . 2008. "Information Ethics: Its Nature and Scope." In *Information Technology and Moral Philosophy*, edited by Jeroen van den Hoven and John Weckert, 40–65. Cambridge University Press.
- Fredrickson, Barbara L. 2004. "The Broaden-and-Build Theory of Positive Emotions." *Philosophical Transactions of the Royal Society B: Biological Sciences* 359 (1449): 1367–78.
 doi:10.1098/rstb.2004.1512.
- Freitas Jr, R. A. 1979. *Xenology: An Introduction to the Scientific Study of Extraterrestrial Life, Intelligence, and Civilization*. Xenology Research Institute (1979-2010).
<http://www.xenology.info/Xeno.htm>.
- Gauvrit, Nicolas, Hector Zenil, Fernando Soler-Toscano, Jean-Paul Delahaye, and Peter Brugger. 2017. "Human Behavioral Complexity Peaks at Age 25." *PLOS Computational Biology* 13 (4): e1005408.
 doi:10.1371/journal.pcbi.1005408.
<http://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1005408>.
- Hammond, Dick. 2005. *Human System from Entropy to Ethics*. Dick E. Hammond.
- Hau, Jorge L, and Bhavik R Bakshi. 2004. "Promise and Problems of Emergy Analysis." *Ecological Modelling*, Through the MACROSCOPE: the legacy of H.T. Odum, 178 (1–2): 215–25.
 doi:10.1016/j.ecolmodel.2003.12.016.
- Heylighen, F. 2007. "Why Is Open Access Development so Successful? Stigmergic Organization and the Economics of Information." In *Open Source Jahrbuch 2007*, edited by B. Lutterbeck, M. Baerwolff & R. A. Gehring, 165–80. Lehmanns Media.
<http://pespmc1.vub.ac.be/Papers/OpenSourceStigmergy.pdf>.
- Heylighen, F., and J. Bernheim. 2000. "Global Progress I: Empirical Evidence for Ongoing Increase in Quality-of-Life." *Journal of Happiness Studies* 1 (3): 323–49. doi:10.1023/A:1010099928894.
<http://pespmc1.vub.ac.be/Papers/ProgressI&II.pdf>.
- Koomey, J., S. Berard, M. Sanchez, and H. Wong. 2011. "Implications of Historical Trends in the Electrical Efficiency of Computing." *IEEE Annals of the History of Computing* 33 (3): 46–54.
 doi:10.1109/MAHC.2010.28.
- Korbitz, Adam. 2010. "Metalaw and SETI: Xenology, Metalaw and Thermoethics." *Metalaw and SETI*. December 2.
<http://web.archive.org/web/20110708053037/http://metalawandseti.blogspot.com/2010/12/xenology-metalaw-and-thermoethics.html>.
- Li, Ming, John Tromp, and Paul Vitányi. 2003. "Sharpening Occam's Razor." *Information Processing Letters* 85 (5): 267–74. doi:10.1016/S0020-0190(02)00427-1.

- Li, Ming, and P. M. B. Vitányi. 1992. "Inductive Reasoning and Kolmogorov Complexity." *Journal of Computer and System Sciences* 44 (2): 343–84. doi:10.1016/0022-0000(92)90026-F.
- . 2008. *An Introduction to Kolmogorov Complexity and Its Applications*. Springer.
- Livio, M. 2000. *The Accelerating Universe*. New York, Wiley & Sons.
- Maienschein, Jane, and Michael Ruse, eds. 1999. *Biology and the Foundation of Ethics*. Cambridge Studies in Philosophy and Biology. Cambridge, UK; New York: Cambridge University Press.
- Mayfield, John. 2013. *The Engine of Complexity: Evolution as Computation*. Columbia University Press.
- Meltzer, Michael. 2010. *When Biospheres Collide: A History of NASA's Planetary Protection Programs*. NASA Aeronautics and Space Administration. https://www.nasa.gov/pdf/607072main_WhenBiospheresCollide-ebook.pdf.
- Moberg, Gary P. 1985. "Influence of Stress on Reproduction: Measure of Well-Being." In *Animal Stress*, 245–67. Springer, New York, NY. doi:10.1007/978-1-4614-7544-6_14.
- Monod, Jacques. 1972. *Chance and Necessity: An Essay on the Natural Philosophy of Modern Biology*. New York: Vintage Books.
- Næss, Arne. 2008. "The Shallow and the Deep, Long-range Ecology Movement. A Summary." *Inquiry*, August. doi:10.1080/00201747308601682.
- Odum, Howard T. 2007. *Environment, Power, and Society for the Twenty-First Century: The Hierarchy of Energy*. New York: Columbia University Press.
- Ostwald, Wilhelm. 1912. *Der Energetische Imperativ*. Leipzig: Akademische verlagsgesellschaft m. b. h. <http://archive.org/details/derenergetische00ostwgoog>.
- Quintelier, Katinka, Linda Van Speybroeck, and Johan Braeckman. 2011. "Normative Ethics Does Not Need a Foundation: It Needs More Science." *Acta Biotheoretica* 59 (1): 29–51. doi:10.1007/s10441-010-9096-7.
- Ramezani, Sara, and Ulle Endriss. 2010. "Nash Social Welfare in Multiagent Resource Allocation." In *Agent-Mediated Electronic Commerce. Designing Trading Strategies and Mechanisms for Electronic Markets*, 117–31. Lecture Notes in Business Information Processing. Springer, Berlin, Heidelberg. doi:10.1007/978-3-642-15117-0_9. https://pure.uva.nl/ws/files/20631648/74912_312938.pdf.
- Regan, Tom. 1987. "The Case for Animal Rights." In *Advances in Animal Welfare Science 1986/87*, 179–89. *Advances in Animal Welfare Science 1986/87*. Springer, Dordrecht. doi:10.1007/978-94-009-3331-6_15.
- Ruse, M. 1985. "Is Rape Wrong on Andromeda? An Introduction to Extraterrestrial Evolution, Science, and Morality." In , 43–78.
- Sidgwick, Henry. 1907. *The Methods of Ethics*. Hackett Publishing.
- Steele, Robert David. 2012. *The Open-Source Everything Manifesto: Transparency, Truth, and Trust*. Berkeley, Calif: Evolver Editions.
- Strick, James Edgar. 2000. *Sparks of Life: Darwinism and the Victorian Debates over Spontaneous Generation*. Cambridge, Mass: Harvard University Press.
- Ulgiate, Sergio, and Mark T. Brown. 2009. "Emergy and Ecosystem Complexity." *Communications in Nonlinear Science and Numerical Simulation* 14 (1): 310–21. doi:10.1016/j.cnsns.2007.05.028.
- Vakoch, Douglas A., ed. 2014. *Extraterrestrial Altruism: Evolution and Ethics in the Cosmos*. The Frontiers Collection. Springer Berlin Heidelberg.
- Varré, J. S., J. P. Delahaye, and E. Rivals. 1999. "Transformation Distances: A Family of Dissimilarity Measures Based on Movements of Segments." *Bioinformatics* 15 (3): 194–202. doi:10.1093/bioinformatics/15.3.194. <http://bioinformatics.oxfordjournals.org/content/15/3/194>.
- Vidal, C. 2014. *The Beginning and the End: The Meaning of Life in a Cosmological Perspective*. New York: Springer. <http://arxiv.org/abs/1301.1648>.
- . 2015. "A Multidimensional Impact Model for the Discovery of Extraterrestrial Life." In *The Impact of Discovering Life Beyond Earth*, edited by Steven J. Dick, 55–75. Doi:10.1017/CBO9781316272480.006. Cambridge: Cambridge University Press. <http://tinyurl.com/vidal2015b>.
- Zenil, Hector, Jean-Paul Delahaye, and Cédric Gaucherel. 2012. "Image Characterization and Classification by Physical Complexity." *Complexity* 17 (3): 26–42. doi:10.1002/cplx.20388. <http://arxiv.org/abs/1006.0051>.