White paper
Scientific Community

Darwinism in the Information Space

How algorithms shape digital territories

Thought Leadership

Executive Summary

The information technology space is, according to Atos' raison d'être, the space where our data and information circulate, where they are stored and processed. Like land, sea, air and space, it is not inert, but conquered and shaped by people[1].

Algorithms are at the core of this shaping. Any digital system is made up of hundreds of thousands of enmeshed algorithms, embodied in operating systems, scripts, applications, business processes, user interfaces, network protocols, APIs, cloud services, machine learning models, digital control loops – everywhere in fact. As these systems become more and more complex, they embed more and more algorithms and execute more and more instructions, to perform never-ending new digital services that deeply shape our digital world. Today's algorithms, in what is called artificial intelligence (AI), are even able to solve the kinds of problems previously reserved for humans, like driving cars or diagnosing diseases. The rate of progress is so high that it raises many concerns about its impact on human societies.

But what are the laws that drive this evolution? How do inhabitants of the information technology space – states, companies, individuals – create these algorithms that shape it? We know that organisms that inhabit the land, air and sea followed a long biological and, in the case of humans, cultural evolution to conquer and shape these spaces. **Could it be the same for the information space?**

This idea appears to be very appealing. In this paper, we will develop it, and explore some facets of how the evolution of algorithms follows a generalized version of the mechanisms of variation and heredity proposed by Charles Darwin.

Author & Acknowledgements

Thierry Caminel,
Distinguished Expert, AI Technology and Innovation Leader

The author would like to thank the Atos Scientific Community, and notably the following members for their valuable comments: Olivier Mallet, Andy Wallace, David Daly, Jose Esteban Lauzan, Peter Venken, Paul Albada Jelgersma, John Hall, Ann Borodina, Erwin Dijkstra, Purshottam Purswani, Dominique Grelet, Benoit Pelletier and Ricky El-Qasem.

[1]Thierry Breton (2019). *Safer, more open and more reliable: the time has come to shape the digital space*

Introduction

Generalized Darwinism, or Universal Darwinism[2], is a generalized version of Darwin's theory of evolution, that avoid some of the confusions related to its biological roots. It has been used first to explain cultural evolution[3], and then applied to a wide variety of other domains like technological innovation[4], anthropology, linguistics, psychology, engineering and cosmology. More generally, it has even been used to understand the emergence of order[5].

According to that approach, many evolutionary processes can be decomposed into just three components: 1/ Variation of a given form or template, typically by mutation or recombination; 2/ Selection of the fittest variants; 3/ Heredity or retention, meaning that the features of the fit variants are retained and passed on.

Reference to Darwinism in this paper should be considered as a shorthand for this process, that considers evolution of the considered entities or patterns (organisms, algorithms, companies, etc.) as a solving procedure to search for the best solution for the problem of how to "survive" (maintain, be retained, be resilient) and "reproduce" (replicate, be copied) by generating new trials, testing how well they perform, eliminating the failures, and retaining the successes.

In the first sections, we will discuss the nature of algorithms, justify our assumption that many human practices can be considered as algorithms, and argue that there's substitutability between these forms and the computer ones. We will then provide examples of Darwinism evolution of digital algorithms in open source communities and in cybersecurity. We will show how AI techniques bring Darwinism at the lower level of code evolution, how they speed up that substitutability, and the crucial role of digital platforms to accelerate the shaping of the digital space.

Putting it all together, we will show some positive feedback loops occurring in the entangled Darwinists processes at several levels of IT, and their similarity with the one in the biological space. That will provide insights enabling us to understand the digital transformation of enterprises and the efforts of companies and governments to control the digital space that, in turn, lead to burdensome considerations about sovereignty in the digital territories, transhumanism and the potential consequences of AI for our societies.

Digital and non-digital algorithms

What is an algorithm? For a computer practitioner, it's basically a sequence of instruction executed by a CPU. But when we explain what an algorithm is to newcomers, we often take the example of cooking recipes. Recipes are a set of instructions that describes how to prepare a culinary dish. Making the dish involves many different algorithms – preparing the ingredients, cooking them, etc. The trick in mastering cuisine is mainly one that involves learning algorithms, that are transmitted by writing, speech, or imitation. On top of this, recipes and cooking culture are good examples of Darwinist evolution: recipes are continuously modified, combined, selected, and transmitted. It's one of human life's pleasures to taste such variations!

Philosopher Daniel Dennett[6] and researcher Richard Dawkins, among others, consider that behaviors and practices spreading from person to person within a culture follow the Universal Darwinism processes. Typical examples include cooking, as we mentioned before, but also music, agriculture, architecture, religion, or war practices – to name a few. All these practices, know-how and behaviors are evolving, possibly by random chance or by mutation and combination of previous ones. The best ones according to some criteria are selected and transmitted through imitation and language.

To emphasize the similarity with genetic evolution, Dawkins coined the term "memes" for these units of a culture or system of behavior that are passed from one individual to another by imitation or other non-genetic means. A few years ago, Yuval Noah Harari went further in his renowned - yet debated - book, *"Homo Deus: A Brief History of Tomorrow"*. One of his keystone ideas is that genetic and cultural practices behave as algorithms. Organisms use biochemical algorithms[7] to "calculate" what is usually considered to be "feelings" or "emotions" — such as the way to make the best decision to avoid a predator or to select a partner for successful mating. These algorithms follow a Darwinist process: They are evolved and improved through millions of years of evolution, and if the "feelings" of some ancient ancestor made a mistake, the genes shaping these "feelings" did not pass to the next generation.

The important point is that Harari argues that digital algorithms and biochemical algorithms have the same nature, and with progress in AI, the former are likely to become more effective for most tasks than the latter. It's one of the key underlying mechanisms that shape the information space, that we will investigate in next sections.

Very recently, other researchers published work on such an extended view of what algorithms are. For example, paleoanthropologist Pascal Picq compared great apes and AI evolution, and coined the term "Darwinian digital space"[8]. Probably more impactfully, MIT Media Lab's researchers published a paper in the journal Nature[9] outlining a broad and ambitious research agenda to study Machine Behavior, in a way that integrates computer science and the sciences that study the behavior of biological agents. One of their goals is to better understand the beneficial or detrimental effects of AI algorithms on society. They notably initiated a framework for studying how algorithms evolve, whose concepts, such as mutation, co-evolution, replication or inheritance, are very close to the one used in that paper.

[2] Wikipedia provides a good overview
[3] Chris Buskes (2013). Darwinism Extended: A Survey of How the Idea of Cultural Evolution Evolved
[4] Andreas Wagner (2014). Spaces of the possible: universal Darwinism and the wall between technological and biological innovation
[5] Campbell, John & Price, Michael. (2019). *Universal Darwinism and the Origins of Order*
[6] Daniel Denett (1995). *Darwin's Dangerous Idea: Evolution and the Meanings of Life*
[7] Yuval Noah Harari (2016). on big data, Google and the end of free will
[8] Pascal Picq (2019). Are machines, like animals, governed by the laws of evolution?
[9] Rahwsan, I., Cebrian, M., Obradovich, N. et al. (2019). Machine behaviour. Nature 568.

Algorithm homomorphism

Let's return to our recipe analogy. The sequence of instruction can be written in a cookbook to be executed by a human. But it could also be encoded with a computer language to be executed in a food production factory; At a given level of detail, the sequence of actions would be similar, just the encoding and the way the actions are performed would be different. We could say that both versions of the algorithms are homomorphic, i.e. they have the same shape. Both can exist and evolve in the digital and in the non-digital information space, and they can move from one space to another.

Image recognition is another example of such homomorphism. It is implemented in organisms by biological algorithms encoded in genes, but also as computer algorithms. They are comparable in the sense that they both have the same functions (classify encoded visual stimuli inputs), and they are potentially exchangeable (a digital vision system can replace a human for many image recognition tasks). This analogy led the robotics expert Gill Pratt to ask the question, *"Is a Cambrian Explosion Coming for Robotics?*[10]", because the "invention" of vision at the Cambrian was the keystone for the explosion of forms of life, and something similar might appear if vision is given to computers. Machines could also learn how the physical world works, for example, by "seeing" thousands of videos, like a child learns about gravity and inertia conservation by observing the world. This is an active and very promising research topic for AI researchers[11].

Darwinist evolution of computer algorithms

Compared to biochemical algorithms, digital ones allow a much faster Darwinist process. The development of the Linux Kernel provides a good example. Many researchers and companies are working to improve it. Whether they work on new concepts or slight algorithm improvements, the path from an improvement idea to the change being introduced into a Kernel release is long, and many ideas or proposals may never make it through. But if it succeeds, then the change will be kept for a long time in the Linux code base, serve as a baseline for future work, and might at some point be deployed onto millions of computers running Linux, from smartphones to supercomputers.

We see that the description of Universal Darwinism above fits nicely with our example:

1/ Variants of the numerous algorithms powering an Operating System are continuously created, typically by changing some part of an existing algorithm or combining different ones (often algorithms developed in another context – something evolutionary biologists call exaptation) ;

2/ The best variations that prove their benefit are selected and put into a Linux Kernel release;

3/ The release is integrated into thousands of products and becomes the basis for new evolutions. In the same way genes survive and continue to evolve after cells die, algorithms in your smartphone continue to evolve after you decide to change it for one with better algorithms (which has, for instance, new or improved features).

From a different perspective, the open-source movement follows such an evolutionary process. On GitHub for example — the largest open-source development platform — millions of projects exist and are continuously forked, mutated and recombined by millions of software developers who create elements of code, or, more commonly, reuse code or templates from other projects. Changes are selected before being introduced into the master branch, some piece of code becomes stable and is reused in many projects or products, while most of them are forgotten. The analogy with genes evolution is sharp: simple organisms are also forked, their genes can be mutated or recombined, and sometimes changes are kept in descendants.

There are other examples of similarities between software and biological evolution. For example, people build phylogenetic trees of Unix variants or programming language. More interestingly, researchers have found that the frequency distribution of usage of crucial bits of genetic code in bacteria is similar to the frequency of installations of Linux packages on computers[12]: both consist of a large number of components (genes or software packages) that can be easily added from a shared pool or removed from individual systems, and the frequency of the most used components follows the same distribution, a power law. Therefore, we can suggest that similar phenomenon of emergence and self-organized criticality occurs in both spaces[13].

[10]Gill A. Pratt (2015). *Is a Cambrian Explosion Coming for Robotics?*
[11]Yann Le Cun (2017). *A Path for AI*
[12]Tin Yau Pang & al (2019). *Universal distribution of component frequencies in biological and technological systems*
[13]Per Bak (1996). *How Nature Works: the science of self-organized criticality*

Darwinism in AI and Machine Learning

Machine Learning (ML) is another way to create digital algorithms, typically by assembling networks of parameterized functional blocks and by training them from examples. The unforeseen progress of the most successful technique today, "Deep Learning", is a good illustration of the Darwinist evolution described above: for decades AI progressed slowly despite its promise and significant R&D funding. A disruption happened around 2011, when incremental amelioration in neural networks training stumbled across the fast performance improvement of Graphical Processing Units (GPU), fueled by the willingness of human gamers to buy ever faster graphic cards. As the matrix algebra to generate graphics is similar to the matrix algebra used to train a neural network, hybridization was possible that led to a spectacular lift for some tasks such as video recognition, and then to coevolution of both deep learning techniques and GPUs. These processors are becoming more and more specialized, leading to an exponential increase of their compute power.

Interestingly, it appears that several of these ML techniques can themselves be considered as using the Darwinist process of variant generation, selection of the fittest, and transmission to a descendant. That is the case for Reinforcement Learning (RL), one of the key ML methods driving the explosion of AI that we are currently seeing. The goal is to optimize a sequence of actions to maximize a goal (for example, in the Go game, a sequence of the placement of stones to maximize the surroundings of the opponent's territory). In that algorithm, an agent (i.e. an algorithm that we train) selects actions by trial and error. It receives a numeric reward, which encodes the success of its outcome, and seeks to learn to select actions that maximize the reward — i.e. the fittest ones. Over time, the agents tend to learn how to maximize the accumulated reward.

Reinforcement Learning is notably used for "Neural Architecture Search"[14] (NAS), a process where an agent automatically generates another neural network for some task, for example to classify images. It's also the core of several automated ML products. The idea is to generate child models in a search space, whose performance is evaluated and used as the measurement of the reward. Many tricks inspired by biological evolution can be used to accelerate the training, such as transferring weights from a successful model to the next generation.

It's also possible to explicitly implement randomness in variant generation in the algorithm itself, like in Genetic Algorithms (GA), a technique invented 50 years ago that has proved to be very effective for solving optimization problems. The idea is to encode solutions in a way they can be combined and mutated, starting from a set of possible solutions (called a population), and then applying the Darwinist process – random mutation, selection of the fittest, transmission to another generation of solutions - until it converges to an optimal one.

Genetic Algorithm techniques are notably used to improve Reinforcement Learning, and that's probably one of the most promising research areas in AI. An early successful technique is "NeuroEvolution of Augmenting Topologies" (NEAT), which uses GA to modify the topology of the Neural Networks, a pattern explicitly based on biological Darwinist evolution principles. Such technique continues to evolve, taking benefit from the progress of gradient-based methods for training deep neural networks[15].

Another approach, also directly inspired by biological evolution, is "Population-Based Training" (PBT)[16], where the concept is to train a population of more or less randomly programmed deep neural networks. The best ones are selected and copied with slight random mutations, and so on for several generations. It was notably used by the Google Deepmind team to develop agents able to win against human professional players in StarCraft II, one of the most complex strategy games, proving that AI systems can simultaneously deal with imperfect information, long-term planning, real-time responses, and a very large action space[17]. Other research from OpenAI has demonstrated that agents in competition or coordination can learn behaviors to solve problems and discover progressively complex tool use, without human supervision or prior knowledge being given[18].

Beyond these examples, all these techniques and their variants are proven to be very effective, and all major AI teams in the world are working on them, either to solve hard AI problems, such as autonomous vehicles or self-learning conversational agents, teaching robots to walk or to automating the process of building machine learning models, without requiring highly skilled data scientists. They are fantastic accelerators to spread and improve AI-powered algorithms everywhere in the digital space.

Darwinism and cybersecurity

Cybersecurity can be seen as the usual battle of the shield and the sword. But as it happens in the digital space, it illustrates well the entanglement between digital and non-digital algorithms. Non-digital algorithms are the security practices or human engineering, whereas digital algorithms are in malware, anti-virus software and in Security Operation Centers. All evolve following Darwinian laws, but the digital ones evolve much faster, and tend to overcome the former.

Computer viruses are a good example, at another level, of the connection between biological evolution and digital algorithms. Their Darwinist evolution has been well studied by researchers[19]: malware modifies the executable code of the target in a similar way that biological viruses modify DNA, they evolve continuously to avoid detection (possibly by mutating their own code to change their signature and behavior) and, to use new exploits, they replicate themselves in the digital information space.

Today, creating such malware requires a lot of human engineering resources, and the most powerful ones are built by nation-controlled agencies[20]. But AI will change this situation. Thanks to evolutionary algorithms, like the one mentioned above that is able to play complex strategic games, it is likely that private groups will be able to build more and more powerful malware.

[14]George Seif (2018). *Everything you need to know about AutoML and Neural Architecture Search*
[15]Matthew Roos (2019). *Evolutionary approaches towards AI: past, present, and future*
[16]Max Jaderberg (2017). *Population based training of neural network*
[17]AlphaStar team (2019). *Mastering the Real-Time Strategy Game StarCraft*
[18]Karen Haon (2019). *AI learned to use tools after nearly 500 million games of hide and seek*
[19]Iliopouloas & all (2011). *Darwin Inside the Machines: Malware Evolution and the Consequences for Computer Security*
[20]Marc Stoecklin (2018). *DeepLocker: How AI Can Power a Stealthy New Breed of Malware*

Darwinism and platforms

Large Digital Platforms, and especially cloud computing platforms — such as the ones owned by Amazon (Amazon Web Services), Microsoft, Google or Alibaba — have a very important role in the evolutionary advancement of algorithms and the shaping of the information space.

Firstly, they are the repository of numerous algorithms, many of them open-source, that they make easily available and composable as frameworks, APIs and highly automated managed services. Digital platforms ease the use of these algorithms for themselves, but also for startups, that can scale up new ideas at low cost by using all these algorithms to develop new applications and services. Some of these startups may, in turn, create new variants of the open source algorithms, which could later be integrated into these platforms. Platforms that succeed to develop de-facto standards - like Google with TensorFlow to develop Deep Learning applications - can increase the retention of new algorithms within their ecosystem, making it more appealing for new developers.

Secondly, these platforms provide services for Machine Learning, that is to create new algorithms from data. These algorithms are used to provide new services that, in turn, contribute to the acquisition of more data, enhancement of algorithms, the gaining of more money to invest in new versions, etc. The largest Cloud companies are fighting to create ecosystems of developers who can create and develop innovative new algorithms, new variants, combine them, etc. By providing the means to host these new algorithms on their platforms, they increase the retention of algorithms in their ecosystem.

Algorithms created from Data can be deployed in many devices at the Edge, and that accelerates the evolutionary improvement process. For example, self-driving algorithms can be created by machine learning techniques from data acquired in real situations. If the algorithm just assists the human (like in the current version of Tesla cars), then it can continuously compare the decision taken by the human driver with the decision it would have taken, and, in the case of a mismatch, or for whatever interesting situation (like a lane change), sends it to the Platform and uses it to train a new version of the machine learning models. This way human driving knowledge is progressively transferred to machines, making them progressively knowledgeable enough to drive a car without human assistance. The rate of progress increases with the number of cars in the field, and each new situation encountered by a car can be used to improve algorithms that can be replicated to other machines via digital networks.

More generally, platforms act as catalysts in the digital information space: by reducing distance and mismatch between them and data, and making the combination of algorithms easier, they enforce the Darwinist evolutionary process by speeding up the creation, combination, selection and the retention of the algorithms. With automated machine learning, these reactions between algorithms in platforms become more and more autocatalytic. Our feeling is that autocatalysis in algorithms is a breakthrough in the shaping of the digital information space.

Feedback loops

As seen before, Machine Learning is a matter of algorithms that can create other algorithms. And they create them mostly from data. These algorithms are used in applications that generate even more data (like user actions, driving situations or tagged images), which in turn allow for the creation of a new generation of (better) algorithms, that can power new services that, in turn, generate more data. It is a positive feedback loop, that happens mostly in the digital platforms we mentioned, that can store in one virtual space data, applications and machine learning algorithms.

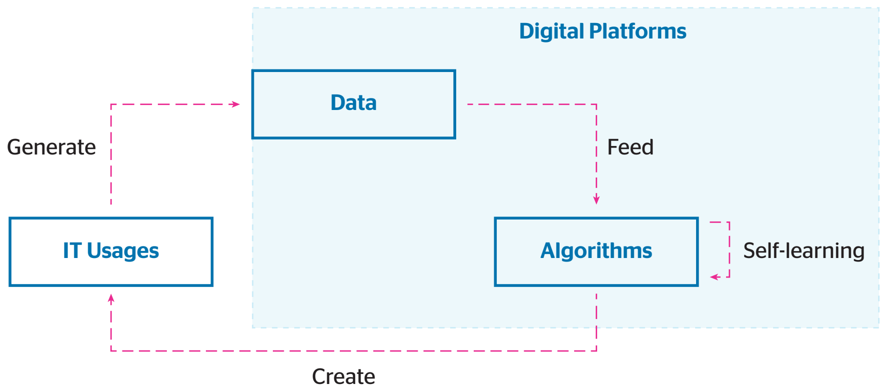

Positive feedback loops enhance or amplify the consequence of an action, and they may cause unpredictability. They can be very important in any evolutionary process because they accelerate it, and they can move a system away from its equilibrium state. They have been used to describe aspects of the dynamics of change in biological evolution. For example, Alfred J. Lotka (1945) argued that the evolution of the species was essentially a matter of selection that fed back energy flows to capture more and more energy for use by living systems.

That does not look very different from what is happening in the digital information space, where data is the energy feeding inhabitant of that space. Companies able to create such a loop can capture a lot of data and thus have large territories in that space, making them extremely powerful. One example, less known but symptomatic, is Blackrock, the world's largest asset management firm. They reached that position because they understood first that algorithms were better than humans for evaluating risks. They are now so large and have so much data and so many algorithms that even governmental central banks have to use their services. In doing so, they give new data to Blackrock, who in turn can create new or better services that can be sold, thus acquiring more new data, etc. And they now use more advanced AI techniques. Although no detail is published, we can figure out that the new evolutionary planification algorithms introduced above will find their place alongside the Monte-Carlo simulations that the firm uses today, allowing Blackrock to provide better investment plans than humans or competitors.

Another example of a positive feedback loop is in the dissemination of knowledge that is accelerated by information technologies — like the Internet, smartphones, on-line courses, AI-enabled teaching and learning, search engines, question-answering systems, etc. These information technologies are backed by many digital algorithms. This means that digital algorithms accelerate the rate of diffusion of cultural algorithms (for example computer science training), which in turn accelerates usage and the development of digital algorithms. It's a form of coevolution, a mechanism known to have an important role in driving major evolutionary transitions.

It's also interesting to consider the recent progress in DNA manipulation, such as CRISPR-Cas9 techniques. Although the information in that space is not digital, this progress is the result of digital algorithmic advances in DNA sequencing and classification. Moreover, DNA can be seen as the physical support of algorithms to create proteins[21].

Therefore, digital algorithms are changing genetic algorithms. It's a feedback from the digital space to the biological one. That could have a tremendous impact in the future, because it will replace the old way of gene evolution through random mutations with one controlled by digital algorithms. These changes are already applied to animals or plants, and many researchers think CRISPR could revolutionize medicine. Transhumanists advocate that the human intellect could be enhanced using such techniques[22], and thus add another feedback loop. The following diagram illustrates these relationships:

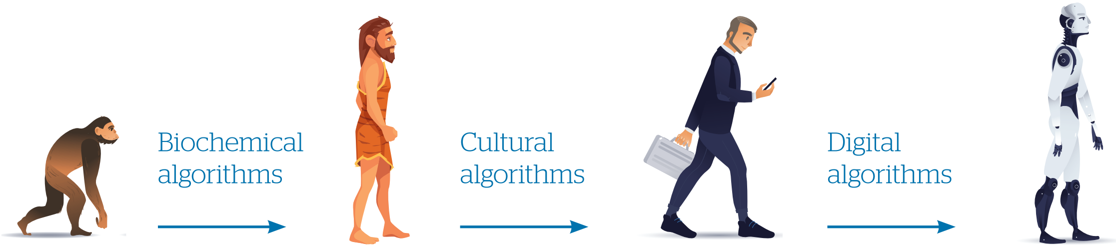

Support	DNA / Genes	Brain / Memes	Digital Storage / Code
Transmission	DNA Replication, biological reproduction	Imitation, language, writing, video, ...	Digital copy, networks, transfer learning (AI), ...
Memorization (examples)	Making of proteins and biological organisms, mating strategy, survival skills	Cultural and agricultural practices, recipes, laws, music, myths, processes, ...	[Subset of previous]
Evolution rate	Millions of years	Centuries	Years
Feedbacks	Modify (transhumanism?)	Speed-up	

[21]Douglas Hofstadter (1982). Metamagical themas, The Genetic Code:Arbitrary?

[22]Antonio Regalado (2019). China's CRISPR twins might have had their brains inadvertently enhanced

Darwinism and digital transformation of enterprises

The digital transformation of processes and competences in enterprises can be considered as a move from human-centered algorithms to digital ones. Sequences of instructions that were performed by humans are replaced by sequences of instructions performed by computers to achieve a similar result. Human know-how previously transmitted by oral or written language is now encoded as computer algorithms in the digital information space. That move by itself is not new: it started at the beginning of automation. The disruption comes from the range of algorithms that computers can now execute, that is increasing exponentially.

There are many, many examples of human tasks that are replaced by computer algorithms. Let's just take one from our industry. In traditional IT systems, humans maintain servers, VMs, databases and applications using a combination of written procedures, simple automated scripts, and know-how acquired by training, imitation or experience. But, when moving to the Cloud, it's "Infrastructure as code". All these tasks are implemented by a myriad of algorithms that automatically add new CPU and storage resources when necessary, configure networks, mitigate failures, replicate and backup databases, analyze logs, prevent faults before they occur, etc. By using these many algorithms implemented as highly automated managed services, a single engineer can build and run in production a system that, otherwise, would require a team of several people. It's a Darwinist process in the sense that different variants of the same algorithm exist – one implemented by humans following procedures and the other implemented in the Cloud by digital algorithms – the one with the best fit (usually which minimizes the money spent) is selected.

This process is accelerated by Machine Learning, which is used to transfer human knowledge to computer algorithms and to increase AI reasoning capabilities.

This competition between humans and machines, as explained in Erik Brynjolfsson and Andrew McAfee's book "Race Against The Machine", is, from an information space perspective, a Darwinist competition between forms of algorithms. Let's illustrate it with another example: the machine-learning-based algorithms from Amazon are now challenging the ones acquired by booksellers to advise a customer. Amazon also develops AI to control datacenters, warehouses, business processes and logistic systems requiring continuously decreasing human intervention. The algorithms behind all this are continuously improved by developers, and many of them are not even employed by Amazon thanks to Open Source communities. They are also continuously improved by using both the huge amount of generated data acquired by the platform and also the progress of Machine Learning. Humans can hardly compete.

Competitiveness and sovereignty in the information space

Algorithms, and particularly those related to AI and cybersecurity, are becoming an essential element of sovereignty for all developed countries, and of competitiveness for companies. The control of the innovation process and platforms where data is stored, and algorithms are shaped, is notably of key importance for controlling the new territories of the digital information space.

The Chinese and US governments for example — and especially military and intelligence-related agencies like DARPA — invest massively in R&D around data processing and AI. The investments are often co-funded by the large commercial platform providers. The outcome is a flow of new ideas and algorithms, and the fittest could at some point be incorporated into these platforms or used to increase nations soft and hard power (like AI-based weapons and Command & Control, Defence intelligence, information warfare, etc.).

Governments can also directly use the services developed by these platform providers, for example, Microsoft who are on the brink of becoming one of the largest US Army contractors. The physical weapons that provide US sovereignty in land, sea and space will be more and more entangled with AI-powered algorithms and meta-algorithms crafted in these platforms.

We see here a clear Universal Darwinist process between states to control the physical space through the digital one. That comes with risks, as was pointed out by Elon Musk who, in reaction to it, funded OpenAI — whose goal is to *"counteract large corporations who may gain too much power by owning super-intelligence systems devoted to profits, as well as governments which may use AI to gain power and even oppress their citizenry"*. The chance of success of such an initiative is debatable.

Conclusion

In this article we have investigated the idea that the evolution of algorithms has a lot of similarities to the biological and cultural evolution that drove breakthroughs in the past, like the Cambrian explosion 500 million years ago or the invention of agriculture 12,000 years ago. Both are based on generating new trials, testing how well they perform, eliminating the failures, and retaining the successes, a process sometime theorized under the name Universal Darwinism or Generalized Darwinism.

We have given examples about how these processes can provide useful insights about how algorithms evolve over time at several scales and in many contexts — ranging from open-source software to platforms, from the digital transformation of our society to the progress of AI systems and the impact on society.

This evolution goes hand in hand with the conquest of the digital information space. It's comparable to the conquest of other spaces, for example the expansion of life in sea from bacteria to humans, that is a continuous Darwinist multi-scale evolution of algorithms at every level of complexity, from DNA mutations in bacteria, to the acquisition of know-how by humans to build ships and organize fleets. Here, the scale goes from mutation in a neural network to enmeshing of thousands of algorithms created by humans or AI. The process is similar.

Although considering social behaviors as algorithms is debated, we think that comparing algorithm evolution independently of their support (genes, culture, or digital) helps to understand the profound digital disruption we are living in. Analyzing the numerous positive feedback loops helps notably to understand the exponential nature of their growth, why their evolution is unpredictable, and the role of large digital platforms as catalyzers. Looking forward, our feeling is that a general theory of algorithm evolution is possible, which could be part of broader studies on machines behavior and, beyond that, the self-creation of order in nature.

There is, however, a fundamental difference between these evolutions: digital algorithms design and selection is driven by humans; they are subject to the scrutiny and ultimate decisions of organizations (private and public), and therefore, can undergo ethical appraisals to guarantee a responsible evolution that is non-damaging to individuals or societies. Positive feedback loops can be analyzed to balance the benign and malign disruptions they can originate. The Darwinist nature of the evolution of the algorithms shaping the digital space can be constrained, for example by preventing them from accessing or use particular information, or by enforcing criteria other than short-term profit for the selection of the "fittest variants". All these controls should be put in place so that organizations could align their purpose with the democratic, ethical and environmental issues to be tackled.

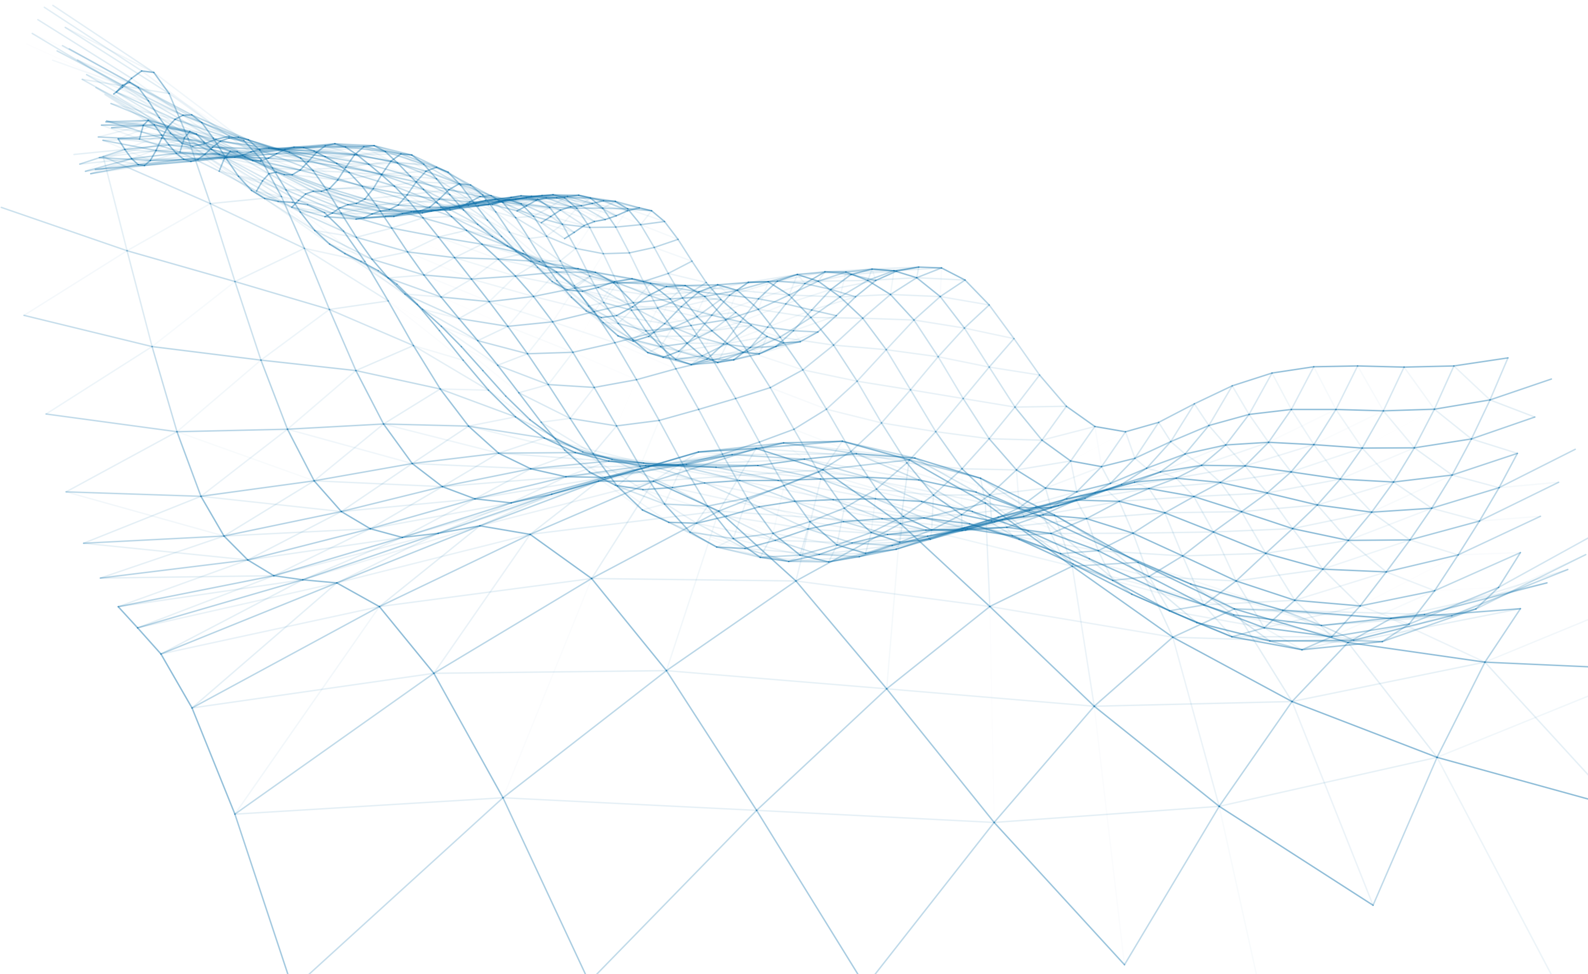

Additional Reading (food for thought)

There are more and more articles on key points addressed in the paper. Here some very interesting ones, published recently and available online:

- Sam Lacey, *Evolutionary Machine Learning - The Next Deep Learning (2019)* - A recent analysis.
- Timothy Prickett Morgan - *Google teaches AI to play the game of chip design (2020)*– Evolutionary AI to optimize AI processors (another feedback loop).
- Cesar Hidalgo, *Why Information Grows: The Evolution of Order, from Atoms to Economies* (2015) –The MIT researcher explains economic complexity and growth from an evolutionary, information self-organization perspective.
- Andreas Wagner - *"From the primordial soup to self-driving cars: standards and their role in natural and technological innovation"* (2016) and *"The Role of Randomness in Darwinian Evolution"* (2012) – Other interesting papers by a renowned biological evolution researcher.
- Meier Lehman, *Programs, life cycles, and laws of software evolution* (1980) – a 40-year-old paper but still very relevant.[abstract]
- K. Johnson, *AI Ethics Is All About Power* (2019) - interesting view on Platforms and States
- François Chollet, *The Measure Of intelligence* (2019) – A profound paper using "algorithm homomorphism" to measure AI, and much more.
- François Chollet, *What worries me about AI* (2018) - short answer: "feedback loops" in social media giving to AI algorithm access to both perception of our mental state, and action over it.
- Michael C. Horowitz *The promise and peril of military applications of AI* (2018)
- George Dyson, *Darwin among the machines, or the origins of artificial life* (1998)
- David Weinberger, *Everyday Chaos: Technology, Complexity, and How We're Thriving in a New World of Possibility* (book, 2019)
- Joel Lehman & al., *The Surprising Creativity of Digital Evolution: A Collection of Anecdotes from the Evolutionary Computation and Artificial Life Research Communities* (2018) *[abstract]*

About Atos

Atos is a global leader in digital transformation with 120,000 employees in 73 countries and annual revenue of € 13 billion.

European number one in Cloud, Cybersecurity and High-Performance Computing, the Group provides end-to-end Orchestrated Hybrid Cloud, Big Data, Business Applications and Digital Workplace solutions through its Digital Transformation Factory, as well as transactional services through Worldline, the European leader in the payment industry. With its cutting-edge technologies and industry knowledge, Atos supports the digital transformation of its clients across all business sectors. The Group is the Worldwide Information Technology Partner for the Olympic & Paralympic Games and operates under the brands Atos, Atos Syntel, Unify and Worldline. Atos is listed on the CAC40 Paris stock index.

Find out more about us
atos.net
atos.net/scientific-community

Let's start a discussion together

CT-200325-RD-Darwinism in the Information Space