

The submerged part of the AI-ceberg

Romain Couillet*, Denis Trystram*, Thierry Ménéssier†

Abstract—This article discusses the contradiction between the current exploding energy demand of artificial intelligence (AI) and of the information and communication (ICT) industry as a whole, and the parallel strong request for energy sobriety imposed by the need to mitigate climate change and the anticipated collapse of civilization as we know it. Under the form of an open reflection on the goods and evils of AI, the article then raises the suggestion of a drastic change in the AI paradigm, more in phase with the vital obligation to design a more resilient society.

I. DEEP LEARNING, THE NEW ELDORADO?

Over the past decade, the considerable growth of the digital world, propelled by artificial intelligence (AI), has had spectacular effects in a few scientific fields such as computer vision or natural language processing, and has given rise to many new technologies and consumer products. Today, this development even claims to revolutionize many other areas of our society. This revolution indeed concerns many aspects of our lives: we (and humanity as a whole) are promised a bright future with more well-being and comfort, a future made of autonomous vehicles, sophisticated human-machine interfaces, humanoid robots for home help, virtual visits to all the museums of the world within reach of a few click, to only name a few [1], [2].

Deep neural network learning is at the forefront of this development and has spread rapidly, far beyond the confidential fields of its beginnings. In a matter of ten years, this specific computer science tool – theorized as early as in the eighties [3] – has reached all levels of the society: in companies, institutions, research laboratories, in virtually all engineering disciplines as well as in life sciences, etc. Easy to use as a black box thanks to an important software development effort – multiple “plug-and-play” solutions have been developed for engineers (and not only computer science experts) such as the popular *tensorflow* library [4], [5] –, deep learning has effectively replaced “conventional” tools (particularly in computer vision and natural language processing), imposing a form of *radical monopoly* on scientific domains.¹

Will deep learning go so far as to replace human beings with brain-like machines to solve all our problems in the same way as for computer-aided vision algorithms, which today “see” objects better than our own brain? Would Asimov’s prophetic cybernetic world really be on the way? [6] Of

course, investing in deep learning and artificial intelligence involves delegating to the machine the power of our decisions, which comes along with many ethics and equity concerns [8]; as Stephen Hawking pessimistically stated it in 2014: “*The development of full artificial intelligence could spell the end of the human race...It would take off on its own, and re-design itself at an ever-increasing rate. Humans, who are limited by slow biological evolution, couldn’t compete, and would be superseded.*” [7]² Yet, the many promises of artificial intelligence clearly tip the scales today towards increasingly more investment in the field [10]. Besides, researchers now deeply investigate the question of fairness in AI to smooth out these thorny angles [9].

II. POSSIBLY, BUT AT WHICH COST?

Consequently, the road is largely open for AI to keep growing and opening new opportunities. This big picture of AI however fails to question the concerns of costs and socio-environmental impact: the actual conception cost of learning algorithms is indeed rarely known by users, or more precisely, as we discuss next, hidden behind the curtain of a massively energy guzzling Internet cloud. More alarming is the production cost of the multiplication of end-user devices resulting from the actual *generation* of new AI-related usage and needs – these new needs (such as autonomous driving, smart cameras) being produced and imposed by R&D scientists rather than requested by society. Here, the “cost” is not to be simplified in terms of numerical performance (complexity, speed, scaling) but rather from an environmental and societal resilience standpoint. And not only the obvious greenhouse gas emissions (63% of the global electricity mix being produced from fossil fuel [11]) [12] but probably most importantly the incurred soil and water pollution, and mineral depletion from metal and fossil fuel mining (80% of the cost of digital equipments over their lifetime is incurred by metal extraction,³ and a modern mobile phone or laptop is composed of more than 50 different metals, many of which being rare and few of which being recyclable [18]), threats on human health and animal lives by the pollution and destruction of their habitat on mining and waste deposit sites, digital fracture between increasingly fewer haves and increasingly more have nots, etc. Figures and details can be found in the reports from the Intergovernmental Panel for Climate (IPCC) [31] or in a growing literature on the topic [14], [37], [36].

*Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG, 38000 Grenoble, France. †Univ. Grenoble Alpes, IPhIG, 38000 Grenoble, France. This work is supported by the “sustainable AI” initiative of the Multi-disciplinary Institute on Artificial Intelligence MIAI at Grenoble Alpes (ANR-19-P3IA-0003)

¹In the sense defined by Illich [33] that it alters the normative system of knowledge generation and sharing: the calls for projects, dedicated conferences, and job offers in data science and deep learning have recently soared and substituted most non-deep learning alternatives.

²As shall be discussed next, this seemingly science-fictional statement is more profoundly explored by Illich [33] on the dangers of societal dependence to oil and machines, induced by an increasing loss of *common knowledge* and *know-how* moved from the population to computers and machines.

³This is often referred to as *embodied* CO₂ emissions. Citing [13], “The embodied energy of the memory chip alone exceeds the energy consumption of a laptop during its life expectancy of 3 years.”

Only recently has this debate on the environmental dangers of AI and ICT been raised [15], [19], [18]. As shown in Figure 1 borrowed from [15], the energy consumption of a single training run of the latest (by 2020) deep neural networks dedicated to natural language processing exceeds 1,000 megawatt-hours (more than a month of computation on today’s most powerful clusters). This corresponds to an electricity bill of more than 100,000 euros (figures in the millions of euros are sometimes found) and 500 tons of CO₂ emissions – that is, the carbon footprint equivalent to 500 transatlantic round trips from Paris to New-York. In comparison, the human brain consumes in a month about 12 kWh,⁴ i.e., a hundred thousand times less, for tasks much more complex than natural language translation. Unlike a mere ten years ago and in spite of the improvement in desktop computer capabilities, it is no longer possible today to train a modern neural network on a personal computer (it would theoretically take up to 405 years according to Figure 1). Yet again, these particularly alarming figures, expressed in cold CO₂ and currency metrics, practically translate in exploding socio-environmental impacts.

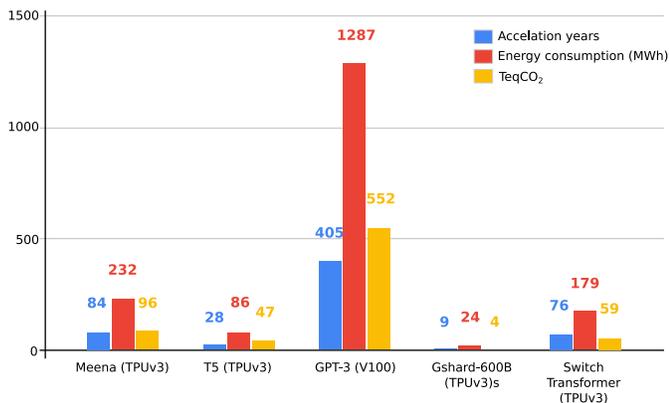


Figure 1. Some factors of energy used to train several modern neural networks (here for language processing); see [15] for detail. Number of GPU-equivalent years (blue), electric bill (in MWh [for information, 1MWh 100€], red), and CO₂-equivalent tonnage (TeqCO₂) [for information, 1TeqCO₂ 1 Return trip Paris-NewYork] (yellow). The mere cost of training a neural network on a targeted application is 250× higher than the annual maximum allowed to each European (2TeqCO₂) to reach the carbon balance in 2050 [16].

One may object that it is probably not surprising that deep learning algorithms be far less energy efficient than three billion years of biological evolution and that the figures may rather suggest a huge room for potential improvement – *neuromorphic AI* technologies being in passing performance-wise and energy-wise quite unsatisfactory at the moment [21]. This objection would displace the focus of the point made here: in a matter of ten years, the absolute consumption of AI learning skyrocketed to reach levels of hundreds of tons of equivalent CO₂ (TeqCO₂) for a single learning task. These levels are at stunning odds with the requirements for the human society to drastically reduce its carbon footprint at a rate of $-7\%/year$, starting today [16]. On the basis of current technologies and the slow development of renewable energy

⁴On the basis of a 2kWh/day consumption for human beings, 20% of which is dedicated to brain activity.

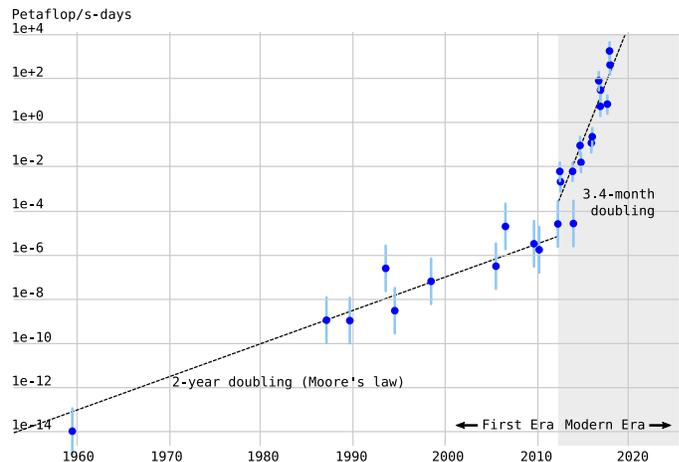


Figure 2. The amount of computation required by neural networks: from Rosenblatt’s perceptron (leftmost) to the latest deep networks (rightmost) [20]. Each blue dot corresponds to an instance of a popular neural network. The growth rate of AI ($\approx \times 11/year$) is a decade higher than the average growth rate of other industrial goods and services, which is already too high to avoid ecological collapse (see next figure).

alternatives (85% of the global energy mix is fossil, a figure unaltered since 30 years [11]), this $-7\%/year$ carbon footprint reduction is strictly equivalent to an *economical degrowth* of $-7\%/year$ starting today.⁵

Let us pursue on the objection line of argument: “once trained”, the algorithm can be reused billions of times, improving billions of users’ satisfaction and well-being, at a comparatively negligible cost on *individual* devices; this would make the initial investment, however large, possibly worth it, if not desirable. The argument here makes the untold assumption that AI algorithms do improve human welfare, which is already a debated position [38], but also hides (i) the fact that practical R&D in AI is a continuous process of trials and errors of deep architectures run on an increasing amount of dedicated servers and (ii) that new algorithms run on up-to-date devices or even dedicated devices, thereby making former equipment obsolete and enforcing the continuous purchase of new terminals. This is the classical *rebound effect*, according to which individual gains in performance or cost (better, less energy-wasteful, algorithms) may indirectly translate into global increase of energetic costs (new equipments, less expensive than more numerous, based on rarer metals, etc.). Rebound effects are particularly marked in ICT, particularly since the emergence of Internet-of-things (IoT) devices [17]⁶: these individual devices tend to increase in number, at a spectacular 67%/year increase rate in the years 2015-2018 [18]. On the server side, the development of deep learning draws behind it an ever-

⁵The calculus is easily made: the IPCC reported in 2011 that a maximum “allowance” of 2,000 gigatons of CO₂ could be release in the atmosphere to prevent a growth in temperature of more than 2°C: at the 40 gigatons/year current rate, and accounting for the loss of ten years of inaction, an exponential decrease in CO₂ emissions not to overtake the 1,0000 gigaton limit needs to operate at a 7%/year rate at least (and 10%/year rate if initiated only in 2025) [16].

⁶In [18], it is reported that out of 57 studied businesses having developed ICT in order to *dematerialize* their activities, *all of them* actually increased rather than decreased their material use (and thus their carbon impact, which we recall relies for 80% on material production).

increasing deployment of specialized computing clusters and network infrastructures that are now hidden in the nebula of the global cloud,⁷ and to which more and more researchers, engineers and data scientists connect on a daily basis.

In a similar manner that the ethical and environmental consequences of animal treatment and meat consumption⁸ is hidden behind the opaque curtain of slaughterhouses to the “grocery-store cellophane-wrapped meat” consumer, this virtualization of computing away from the modern engineer’s laptop additionally induces a new illusion of “harmlessness” of AI and of the digital society. In reality, the manufacture and use of the Internet cloud – which is anything but immaterial – generates hundreds of tons of additional CO₂ emissions (from metal extraction, some of which rare, production, transportation and usage) and is sometimes even accompanied by the co-building of coal-fired power plants to support the intense electrical demand for operation [23]. This growth in connected devices production and server development contributes to the prediction of an increase of the global share of electricity consumption for digital services from 3-4% today to 20% in 2050 [18].

This staggering increase of CO₂ emissions from device production consequently spawns more electronic waste, mostly non-recyclable, which ought to be treated at the end of the equipment’s life. In reality, only 17% of waste are *legally* recycled, this being due both to modern equipments containing compounds of tens of different metals that have become impossible to disassemble, to the comparatively important cost of recycling versus mining [18], [24]⁹, as well as to the fact that many devices are either not disposed of (they stay in our drawers) or trashed as domestic waste. Moreover, even if an unrealistic 100% of metals could be recycled at no cost, intensive mining of new material would still be required to cope with the ever-increasing demand of new material products. The *externalities* of ICT material production – that is, the indirect costs not accounted for at production time – are numerous and mostly unfathomable to the common end-user: irreversible soil, aquifer and river pollution, deforestation and habitat destruction, on and far around mining sites; increase of CO₂ production for mining, production, transport, and electricity consumption for usage of every produced device; deteriorated health conditions of increasingly more human workers on mining sites; worldwide biosphere depletion due to pollutants breaking the trophic chains, etc.

III. SO DOES AI ENDANGER THE PLANET

⁷Yet, comparatively, the volume of shared servers has a lesser impact than individual billion-fold multiplied individual devices.

⁸The effective number of live animals, fish, bird, or Earth mammals, slaughtered annually for human consumption is evaluated to 300-400 billions (approximately one animal per week per person), among which more than 70 billions land animals. This annual figure corresponds to three-to-four times the number of human beings that ever existed on the planet [22].

⁹An interesting mental image (borrowed to V. Mignerot) is that of a sand bucket extracted from a sand beach and then spread out on the parking lot: inexpensive to extract from the “mining site”, the energy needed to collect the same amount of sand when spread out is comparatively huge.

AND OUR CIVILIZATION?

To what extent can humanity bear these multi-spectral costs? A long series of studies, regularly synthesized in the reports from the Intergovernmental Panel for Climate Change (IPCC) [26], confirm that the volume of CO₂ releases into the atmosphere due to human activities over the past fifty years threatens the stability of the Earth system, of life on the planet, and of the human civilization as a consequence [37], [14]. In the so-called *business-as-usual* scenario, where humanity maintains a heavy foot on the CO₂ gas pedal towards a wall that is approaching at increasingly high speed, these studies envision a collapse of the whole human civilization and of life on Earth between 2025 and 2050, as a consequence of having overtaken several no-return thresholds (such as mass extinctions of species,¹⁰ the appearance of now uninhabitable hot and humid zones,¹¹ spontaneous forest fires, total and irreversible erosion of soils, emptied fossil aquifers, etc. [14], [27]). This already unfolding societal collapse¹² goes hand in hand with a drastic reduction in energy resources that cannot be replaced by renewable sources – at least at the level of our current and inexorably growing consumption [25], [16]. This is particularly true of oil, whose extraction peak is imminent [28]. Additionally accounting for the fact that *immediate* measures must be carried out with no further delay and thus with existing technologies, the blind technoptimistic approach consisting in waiting for “*the* technology that will save the planet” (with less energy and shrinkingly less time) is unreasonable – all the more so when approaching the question from a thermodynamic standpoint. In addition, figures show that the effective efforts towards a decarbonated economy of the past decades only led to a continuous increase in global CO₂ production [29]¹³ and that there is no sign in sight of a disrupting evolution in energy production.

The above statement of an imminent collapse of the western civilization is forcefully supported by mathematical models. Figure 3 displays the said *business-as-usual* scenario from the popular “The Limits to Growth” report [30] in its last 2005 update. The initial predictions of the report (as early as 1972) are still valid today and foresee a collapse of the society by 2030.¹⁴ For many authors and research groups (the IPCC [31] worldwide, the Shift Project [16] or the Negawatt initiative [25] in France, the Drawdown Project [32] or Plan-B [14] in the US, etc.), the only viable response to ecological

¹⁰Although largely debated, the rate of life extinction is estimated to occur at a 10 to 100-fold faster rate than during the Cretaceous extinction (that saw the extinction of dinosaurs), and 1,000 to 10,000-fold faster rate than *background level* (that is, “natural” extinction rates).

¹¹Human beings cannot survive more than a few days under a 35°C temperature with 100% humidity.

¹²In south-hemisphere countries notably: see in [14] for instance a report on the increasing number of *failing* states (that is countries no longer in capacity to support their own subsistence).

¹³Only disrupted by the Covid-19 crisis but immediately recovered the year after.

¹⁴This “doom date” needs not be taken as inevitable nor predictive: close to the tipping point of collapse, policies will emerge that will inevitably affect the shape of the anticipated trajectory by changing the equations in the model. Yet, the qualitative nature of the results have been verified not to be altered if no drastic and *systemic* (i.e., of the whole thermo-industrial economy) political shift is performed.

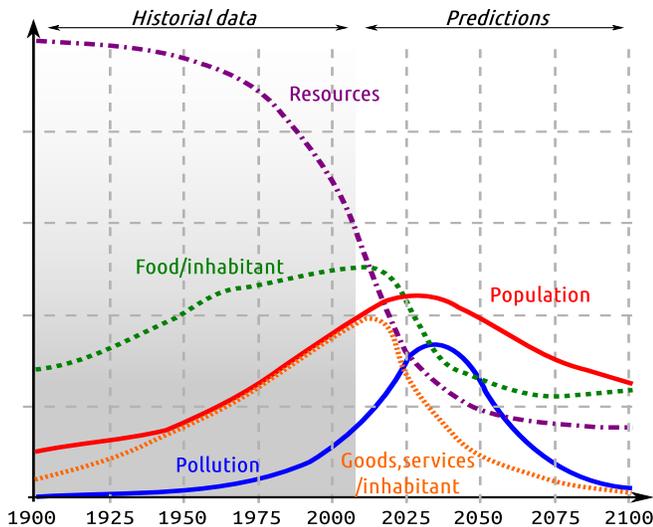


Figure 3. The “World3” model in the *business as usual* scenario [30]. Created in 1972, these predictions are still valid in 2021 – and their robustness has been the subject of multiple validations. The model predicts a population collapse induced, in order, by (i) the increase in pollution that has become economically impossible to curb and which, (ii) due to a lack of resources, dominates the production of goods and services, and then (iii) of food production (itself strongly impacted by pollution and soil erosion). The lack of equipment and famines then (iv) cause mortality to rise and the population to collapse.

collapse is to commit our western civilization to a drastic reduction in the production of goods and services drawn directly or indirectly from the planet (minerals, energy, water), to reconnect human beings (reduce inequalities) and other-than-humans (reinstating *nature* rather than *Homo Sapiens* at the center of concerns), and to rethink the symbolism of *growth* [33], [34], [35], [37], [36]. Dodging the imminent collapse would be a matter of aiming for a growth in *collective well-being* (humans and other-than-humans) which would support a necessary *degrowth of the thermo-industrial economy*. As recalled earlier, the rate of necessary degrowth to fulfill the Paris Agreement (of maintaining global warming below $+2^{\circ}\text{C}$) is evaluated, based on data updated to 2021, at least to 7% per year if it were initiated today, and to even more if nothing is done in the coming years (until it becomes impossible to avoid the complete collapse of the society) [16]. In concrete terms, this amounts to dividing the current energy consumption of the average European citizen by 5 by 2050 and of the average American citizen by 10.

This naturally raises the question of what should be eliminated or reduced from the industrial pressure on the planet, as well as socially acceptable strategies for achieving the objectives of degrowth: which tools can we afford to keep, which technologies should be absolutely abandoned? Knowing that we must aim for an 80% cut (in Europe) of all our daily activities in a matter of 30 years.

Applied to the seemingly marginal field of artificial intelligence and deep learning, which we however showed to develop at an exploding pace and to trigger rebound effects through material production, this pressing need to rethink our priorities *fast* leads us to reflect on the very position of research and investment in AI, and dare ask in full honesty whether AI

should be abandoned altogether.

IV. SHOULD WE CLOSE THE TAP OF RESEARCH AND INVESTMENT IN AI?

Artificial intelligence, and particularly deep learning, have drawn public attention to the point of reaching fields which had never “felt” the need for artificial intelligence but which now seem to consider it urgent not to lag behind (domains as diverse as finance, medicine, biology, environment, etc.). Signal processing and its various branches are no exception (even wireless communications, despite the optimality of Shannon’s theory and a long history of mathematically solid research, now propose encoding-decoding strategies based on deep learning black boxes), with an increasingly growing share of articles on deep learning applications submitted to leading conferences, such as ICASSP or SSP, and journals (Transactions on Signal Processing) of the community. In an imminent world of severe energy constraints and unsustainably high pollution levels, it would come to reason that researchers rather heavily invest on “tools for resilience”: that is, technologies which rely on little energy expenditure, do not require the extraction of further raw (especially rare) material, and whose usage has a positive impact on the planet and humanity’s well-being, and which target resilience to a contrived and unsettled environment. Illich [33] back in the seventies, or Stiegler [34] more recently, provide a clear picture of what ought to be a *convivial* (that is, resilient¹⁵) technology. These are worlds apart from the path followed by researchers today. What is indeed the added value of generating new technologies and algorithms, which may be potentially better in terms of “classical” performance indicators (data rate, error rate, carbon footprint), but which are comparatively more harmful for the planet, and thus clearly work *against* resilience? The whole question precisely lies in the word “comparatively”, which opens a discussion on trade-offs. Comparing the relevance of tools and technologies that produced social progress is not as measurable a task as evaluating an error rate or a probability of detection, especially as it involves non-technical parameters of a social, ethical or psychological nature. But, as pointed out earlier [16], the more we wait to initiate the necessary energetic degrowth without adapting our technologies, the steeper the degrowth slope will have to be operated once started: this will translate in purely abandoning more technologies.

Which brings us to the question of “usage”: the bulk of the carbon footprint due to AI is obviously not confined to our scientific labs, but rather to the leisure applications and comfort devices that the industrial society creates every day and that very quickly grow from the status of gadgets to that of indispensable needs.¹⁶ In addition to harming the planet, the very fact that technologies reinforce the societal *dependence*

¹⁵The notion of *conviviality* is more profound than that of *resilience*, the latter being a mere consequence of the former: a convivial technology also accounts for the level of dependence it induces on its users, for the alteration of their lifestyles (*radical monopoly*), for the consequences of its normative power (*homogenization*), etc.

¹⁶Spam detection, automated suggestion, or face recognition appear as the bulk applications of AI, none of which can be claimed to bring socio-environmental advantages.

on energy and the illusion of “plentifulness” further reduce our degree of resilience and consciousness of a fateful future: as the classical image goes, we are driving fast but so far mostly unharmed on an increasingly shaky road and our best answer is to further accelerate and loosen the seatbelt. Paradoxically though, the first blow to a technological dismantling could actually be an easy one: does it not stand as obvious, if not a moral duty when heading for a renewal of our life symbols (aiming for the protection of living beings, humanity included), to initiate the *degrowth of AI* by giving up on most of these gadgets and then concentrate our efforts on *a priori* more virtuous uses?

But what are these virtuous deep learning or AI-based technologies worth investing in, if any? In anticipation of energy becoming scarcer, the technical tools that are expensive to study *for themselves* are already unsustainable. In the context of deep learning, one can already anticipate an upcoming tipping point where the costs of the very research in modern AI will take precedence over the performance gains, however significant, brought about by neural networks. This is well documented in [19] which shows that deep learning has reached a plateau and that further research in the development of deep learning as it stands is unsustainable.¹⁷

Does this mean the end of artificial intelligence altogether? Not necessarily. But perhaps should we approach AI with a fresh look. An eloquent example is provided by the work of d’Acremont, who concretely demonstrates the value, in terms of both resilience *and* performance, of a “low-tech” AI technology which does not even use a computer in the end. Her application targets the diagnosis of sick children in Tanzanian dispensaries, in the absence of professional doctors, but also under frequent losses in electricity to run computer-dependent algorithms (and accounting for the consequences of hardware failure) [40], [41]. As d’Acremont perfectly illustrates, the challenge here lies in our ability to extract sufficient interpretability from artificial intelligence machines to *then* instruct the dispensary workers (the AI end users) of the lessons learnt from the machine while ultimately dispensing with the machine itself – what ultimately remains in the end is a handwritten decision tree and *improved human knowledge and know-how*. This is an extremely resilient approach: for once, it has the effect of ultimately placing the knowledge produced by AI back into the hands of the human experimenter, a fundamental feature of conviviality; also, it prepares the human experimenter to the moment when, sooner or later, one may be forced to abandon the AI tool permanently – in the context of the unreliable Tanzanian electricity grid, this “moment” has already come.

¹⁷Precisely, [19] shows on empirical data from the past ten years that halving the error rates in deep learners requires a multiplication by a factor 2^9 of the effective number of operations (flops) during training: further improving deep learning thus comes at an exponentially prohibitive cost. The latest architectures in date being trained for weeks or months on thousands to tens of thousands of parallel GPU cores, the upper boundary of deep learning performance has virtually been reached.

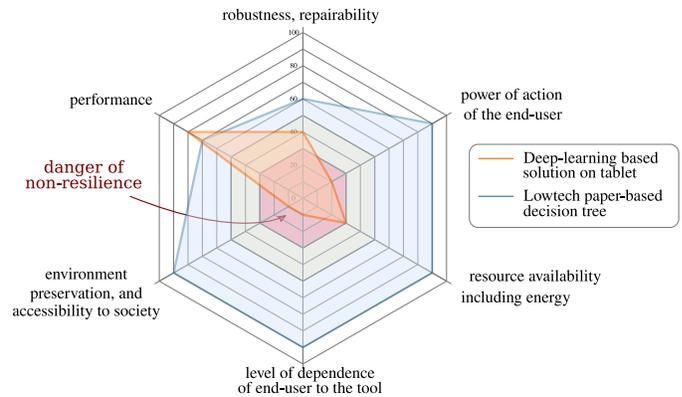


Figure 4. Confronting high-tech and low-tech solutions in terms of *resilience*. Here in the context of d’Acremont two solutions to children diagnosing in Tanzania: high-tech black-box solution on tablet (orange) versus low-tech decision-tree on paper (blue) [40], [41]. Almost no resilience criterion is met by the high-tech solution (unusable under electricity losses, vital risks in hardware failure, black-box effect reducing self-knowledge but increasing technology-dependence of the end-user, socio-environmental impact of building new tablets if solution is spread across Africa, etc.).

V. PROMOTING LOW-TECH LEARNING

D’Acremont’s telling example is not isolated. They do not make the headlines today, but we are already observing the emergence of basic but efficient alternatives to modern AI along with a realization that AI is not everywhere indispensable. Going further, an increasing interest is laid on *low-tech* tools [45], [46] and in the development of “low-tech labs”, for the moment not within universities *per se* but in civic associations [47]. The concept of low-tech fundamentally relies on the design of technologies based on a minimalistic use of resources (reuse or recycling of existing material and equipment, extraction of little and renewable energy), on simplicity and durability (reparable by anyone with basic knowledge and basic material), limited in cost (accessible to a vast majority of people), and adapted to local resilience (in time and space). These precepts of the low-tech movement – minimal cost in local resources, degree of simplicity, value of usage – could naturally define a *new metric*, call it “technical conviviality”, to aim for in future scientific communications, in place of the unsustainable metric of absolute performance (and even of the overly restrictive CO₂ impact). A particularly unconventional – but at the same time crucial – aspect of this *new metric* lies in its revolving around *non easily measurable* defining features of conviviality (level of dependence or power of action of the end-user, reparability of the tool, etc.). It is fundamental that these difficulties do not serve as a pretext to stick to a “classical” reassuringly fully-measurable approach to design a resilient signal processing – which is unfortunately the road taken by carbon code analyzers, green IT or green AI proposals. Figure IV provides a typical representation (with purely illustrative values) of the differing levels of technical conviviality obtained by high-tech versus low-tech alternatives in the context of d’Acremont’s work, underlining the limitations of high-tech solutions which crosses several red lines of non-conviviality.

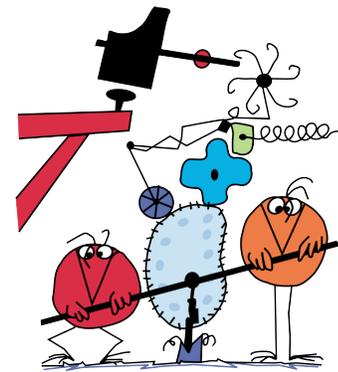
In the same line of thoughts, but even more profoundly rooted in resilience, is the incentive for a shift from the present technology-based economical growth in direction of a *symbiotic* (or *positive*) economy [42], [43]: this radically new approach relies on the capability of *regeneration* of resources by nature and proposes to reinstate natural material and processes as a replacement for rare metals and industrial technologies. A large-scale and long-term symbiotic economy “chance experiment” was run in Portland as a consequence of introducing vegetal alternatives to (unaffordably expensive) wastewater treatment facilities, with positive snowballing consequences on inhabitant well-being and lifestyles (more vegetal in town, lower temperature in summer, less cars and more bikes, improved local and organic agriculture, reduced levels in violence, etc.) [44].

These ideas of sharing knowledge and simplifying the industrial process to focus on convivial tools fully embrace the request for urgent measures of fast industrial degrowth and of an increase in technical resilience (along with an increase in well-being induced by a reconnection to nature and meaning). As the Shadoks put it in Figure 5, in a spirit closer to the early days of signal processing (when energy, memory and device capabilities were more constrained and solutions less wasteful), “why do simple when you can do complicated?” Yet, despite the obvious (there does exist a wide range of solutions to mitigate the impact of the thermo-industrial society on the planet [48]), these very ideas remain hard to promote in academic conferences or calls for projects, which lag behind rather than ahead of both society’s demands and needs.

In the present case of research in AI, turning a blind eye on the industrial fundings in deep networks research is a challenge in itself. In our present *narrative* of a shining future career, it is conceived difficult for a young researcher or engineer to avoid the call of deep learning, let alone AI altogether. This observation is all the more alarming that these young talents are being labelled as the saviors of a planet that their daily work continuously degrades. The resulting strong *cognitive dissonance* is increasingly taking the form of new psychological diseases of civilization such as eco-anxiety and solastalgia, which are now filling the offices of clinical psychologists with patients of a type never seen before [49] and push graduate students from prestigious schools and universities to refuse being recruited in multinational companies [50]. But is this image of a fruitful career in the already overcrowded and essentially mathematically-stuck field of AI so genuine? After all, who may claim authority today in the so far essentially empty domain of *technological resilience*, whose stakes will rise fast? Could the early days of a formerly *resource-constrained* signal processing, rooted in powerful yet accessible mathematics and composed of a myriad of now-forgotten smart and efficient tricks, illuminate the path? [51]

VI. CONCLUSION

To state it bluntly, if it is agreed that the purpose of research in AI, along with any other form of research, is to improve



POURQUOI FAIRE SIMPLE QUAND
ON PEUT FAIRE COMPLIQUÉ ?!

Figure 5. The Shadoks: “why do simple when you can do complicated?”, image reproduced from <http://www.lesshadoks.com/index2.php?page=3>. The Shadoks are fictional bird-like cartoon figures who live in a strange world which they try to harness by elaborating complex never-working machines. They metaphorically illustrate our society made of blindly complex and sometimes useless technologies.

knowledge in order to maximize collective well-being, why are we deliberately jeopardizing the survival of humanity and annihilating the commitment of our young talents when more virtuous and less costly alternatives are already within reach? Perhaps should we feel enthusiastic about this unique opportunity to contribute to a necessary scientific renewal and societal improvement in conviviality, before we simply run out of time?

REFERENCES

- [1] Chakraborty, C. (2021). Artificial intelligence and the fourth industrial revolution. Jenny Stanford Publishing.
- [2] Daugherty, P. R., & Wilson, H. J. (2018). Human+ machine: Reimagining work in the age of AI. Harvard Business Press.
- [3] Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feed-forward networks are universal approximators. *Neural networks*, 2(5), 359-366.
- [4] Masood, A., & Hashmi, A. (2019). Cognitive Computing Recipes: Artificial Intelligence Solutions Using Microsoft Cognitive Services and TensorFlow. Apress.
- [5] Ravichandiran, S. (2018). Hands-on reinforcement learning with Python: master reinforcement and deep reinforcement learning using OpenAI gym and tensorflow. Packt Publishing Ltd.
- [6] Asimov, I. (1940). I, Robot. A collection of short stories.
- [7] Delipetrev, B., Tsinaraki, C., & Kostic, U. (2020). Historical Evolution of Artificial Intelligence.
- [8] Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6), 1-35.
- [9] Bird, S., Dudík, M., Edgar, R., Horn, B., Lutz, R., Milan, V., ... & Walker, K. (2020). Fairlearn: A toolkit for assessing and improving fairness in AI. Microsoft, Tech. Rep. MSR-TR-2020-32.
- [10] Corea, F. (2017). Artificial intelligence and exponential technologies: Business models evolution and new investment opportunities. Springer.
- [11] Electricity mix (2021). Our World in Data. <https://ourworldindata.org/electricity-mix>
- [12] Hoekstra, A. Y., & Wiedmann, T. O. (2014). Humanity’s unsustainable environmental footprint. *Science*, 344(6188), 1114-1117.
- [13] Embodied energy of digital technology (2009). LowTech-Magazine. <https://solar.lowtechmagazine.com/2009/06/embodied-energy-of-digital-technology.html>
- [14] Brown, L. R. (2008). Plan B 3.0: Mobilizing to save civilization (substantially revised). WW Norton & Company.
- [15] Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L. M., Rothchild, D., et al. (2021). Carbon emissions and large neural network training. arXiv preprint arXiv:2104.10350.

- [16] The Shift Project (2017). Modeling global emissions trajectories compatible with a +2°C carbon budget.
- [17] Mordor Intelligence (2021). Internet of Things (IoT) Market - Growth, Trends, COVID-19 Impact, and Forecasts (2021-2026)
- [18] The Shift Project (2019). Lean ICT – Towards Digital Sobriety. Report of the working group directed by Hugues Ferreboeuf for the think tank the Shift Project.
- [19] Thompson, N. C., Greenewald, K., Lee, K., & Manso, G. F. (2020). The computational limits of deep learning. arXiv preprint arXiv:2007.05558.
- [20] AI and compute (2021). OpenAI blog. <https://openai.com/blog/ai-and-compute/>
- [21] Miranda, E., & Suñé, J. (2020). Memristors for neuromorphic circuits and artificial intelligence applications. *Materials*, 13(4), 938.
- [22] Animals slaughtered for meat (2020). Our World in Data. <https://ourworldindata.org/grapher/animals-slaughtered-for-meat>
- [23] Glanz, J. (2012). Power, pollution and the internet. *The New York Times*, 22.
- [24] Electronic Waste. Wikipedia. https://en.wikipedia.org/wiki/Electronic_waste
- [25] Negawatt association (2017). NégaWatt scenario 2017-2050. https://negawatt.org/IMG/pdf/181128_negawatt-scenarioi_eng_12p.pdf
- [26] Bernstein, L., Bosch, P., Canziani, O., Chen, Z., Christ, R., & Riahi, K. (2008). IPCC, 2007: climate change 2007: synthesis report.
- [27] Effects of Global Warming by Region. IPCC. <https://timeforchange.org/effects-of-global-warming-by-region-ippcc/>
- [28] Jancovici J.-M. (2015). Peak oil? Did it already happen somewhere? <https://jancovici.com/en/energy-transition/oil/peak-oil-did-it-already-happen-somewhere/>
- [29] CO2 emissions (2020). Our World in Data. <https://ourworldindata.org/co2-emissions>
- [30] Meadows, D., & Randers, J. (2012). *The limits to growth: the 30-year update*. Routledge.
- [31] Intergovernmental panel on planet change (IPCC). <https://www.ippcc.ch/>
- [32] The Drawdown projet. <http://www.drawdown.org>
- [33] Illich, I. (1973). *Tools for conviviality*.
- [34] Stiegler, B. (2010). For a new critique of political economy. *Polity*.
- [35] Barrau, A. (2019). Le plus grand défi de l'histoire de l'humanité – Face à la catastrophe écologique et sociale. Michel Lafont. In French.
- [50] Gosset A. (2021). Ruptures, le film. (French movie) <https://www.ruptures-le-film.fr>
- [36] Servigne, P., Stevens, R., & Chapelle, G. (2020). *Another End of the World is Possible: Living the collapse (and not merely surviving it)*. John Wiley & Sons.
- [37] Servigne, P., & Stevens, R. (2020). *How everything can collapse: a manual for our times*. John Wiley & Sons.
- [38] Beinsteiner, A. (2020). Conviviality, the Internet, and AI. Ivan Illich, Bernard Stiegler, and the Question Concerning Information-technological Self-limitation. *Open Cultural Studies*, 4(1), 131-142.
- [39] Naess, A. (2009). *The ecology of wisdom: Writings by Arne Naess*. Catapult.
- [40] Rambaud-Althaus, C., Shao, A., Samaka, J., Swai, N., Perri, S., Kahama-Marro, J., ... & Genton, B. (2017). Performance of health workers using an electronic algorithm for the management of childhood illness in Tanzania: a pilot implementation study. *The American journal of tropical medicine and hygiene*, 96(1), 249.
- [41] D'Acremont V. (2021). Technologies et santé : Quels compromis entre éthique, environnement et climat ? Analyse réflexive et expérience de terrain. https://www.youtube.com/watch?v=oKcy_cY0QOw (in French)
- [42] Sury, B., Kestemont, M. P., & Lejeune, C. *Positive Economy, its Concepts and Academic Integration*.
- [43] Delannoy, I. (2017). *L'économie symbiotique: Régénérer la planète, l'économie, la société*. Éditions Actes Sud. In French.
- [44] Abbott, C. (2004). *Centers and edges: Reshaping downtown Portland. The Portland edge: Challenges and successes in growing communities*, 164-183.
- [45] Bihouix, P. (2014). *L'Âge des low tech. Vers une civilisation techniquement soutenable: Vers une civilisation techniquement soutenable*. Média Diffusion. In French.
- [46] Radjou, N., Prabhu, J., Ahuja, S. (2015), *Frugal Innovation: How to do better with less*, Economist Books.
- [47] Low Tech Labs. <https://lowtechlab.org/en>
- [48] LowTech Magazine (under its own low-tech implementation). <https://solar.lowtechmagazine.com/>
- [49] Ojala, M. (2018). Eco-anxiety. *Rsa Journal*, 164 (5576), 10-15.
- [51] Adali T., Trussell H. J., Hansen L. K., & Calhoun V. D. (2018). The dangers of following trends in research: Sparsity and other examples of hammers in search of nails. *Proceedings of the IEEE*, 106 (6), 1014–1018.