Conceptualizing future scenarios of artificial intelligence: from energy servants to AI servants

Richard A Roehrl (DESA) [1]

I suggest taking a big picture view takes a historical perspective and one grounded in energy and thermodynamic considerations. Why? Human history is essentially a story about energy conversions. Energy conversion drives all changes. No structures are possible without energy inputs. Above all, no life is possible without energy inputs. It is fundamental to food and water. It is embodied in technologies, buildings, and infrastructures. Energy is the master resource - the ultimate resource. From a thermodynamic point of view, computation is closely related to the concepts of entropy and energy.

Learning from the past – how we overcame the limits of human and animal muscle power

For most of human history, muscle power was the limiting factor. Our human bodies are essentially energy conversion machines. They convert food energy (based on photosynthesis running on solar radiation) into work, thermal energy, chemical energy in fatty tissues, and electric potentials in our brains. Roughly three quarters of calories burnt go into sustaining basic life functions *at rest.*[2] This is the so-called basic metabolic rate of roughly 83W at rest (e.g., while sleeping). When we do useful work, our body's energy used per time (power) will temporarily be much higher[3]. The total power input for a human differs by age and gender, but averages out to about 113 W. The average daily work-energy requirement in the form of food for a very heavy workload is 2.9kWh, corresponding to an average power input of 120W. Hence, a human who is physically hardworking 365 days a year needs a food energy intake of about 3.8GJ per year.

The industrial revolution led to modern energy technology systems that arguably have freed billions of people from physical toil and slavery. Modern energy fuels can do the job much more effectively and at a much larger scale. Indeed, by 2019, the average world primary energy use about 81 GJ per person. This is roughly 21 times the food energy intake of a physically hardworking man. In other words**,** on average everyone of the 7.7 billion people on our planet – babies and old people included – had available the equivalent of 21 "energy servants" – modern energy converters that can do the physical work of 21 strong men. Instead of "energy servants" we could also call them "artificial muscles". This illustrates the scientific and technological prowess that is behind today's material standards of living which – for most – are much higher than in the past. On the other hand, these 164 billion "energy servants" also mean that our environmental and economic footprint maybe more like 170 billion people (equivalents) rather than 7.7 billion. In fact, the number of our "energy servants" has grown much faster than human population, which is primarily responsible for our ever-increasing global environmental footprint.

Table 1. Number of deployed "energy servants" providing physical labour, 1800-2019

	1800	1900	1970	1990	2000	2010	2019
World population [billion]	1.0	1.6	3.7	5.3	6.1	7.0	7.7
"Energy servants" (human muscle power equivalents) [billion]	5.3	11.5	63	100	115	144	164
"Energy servants" per person	5	7	17	19	19	21	21

Now the analogy between the physical machines and humans is, of course, far from perfect. To compare impacts, our primary energy method is ideal, but to see how machines have provided artificial muscle power to complement and replace human muscle power, it would be better to compare useful energy for which unfortunately global data is unreliable and sparse. Also, an increasing share of primary energy is used for computation and data writing including for running DNNs and other AI algorithms – issues that we will

[1] Note: The views expressed in this brief are those of the author and do not necessarily reflect those of the United Nations or its senior management.

[2] E.g., liver, spleen, brain, skeletal muscle, kidney, heart, etc.

[3] from 125W while standing, 280W while walking, 740W running cross-country and 2415W sprinting

discuss later. However, these energy demand categories are still comparatively small. In addition, energy efficiencies of the early steam engines started very low and lots of efficiency improvements have been achieved especially in the last few decades. However, energy system efficiencies have not changed that much as higher quality fuels and services, such as electricity, have been adopted. In any case, our method provides a simple, order of magnitude illustration of the large global sustainability and distributionary impacts of an ever-increasing system of machines performing physical labour for us. Similarly, it can be used to illustrate the large technological and socio-economic divides between countries and population groups within countries.

Figure 1 shows the number of implied "energy servants" per person in the world. People in some rich countries command several hundred "energy servants" each. People in the poorest countries command less than one each. For example, primary energy consumption in Singapore in 2019 amounted to 611 GJ per capita. Hence, every Singaporean had - on average - the equivalent of the muscle power of 160 men at their disposal.

Figure 1. "Energy servants" per capita providing physical labour in countries around the world in 2019

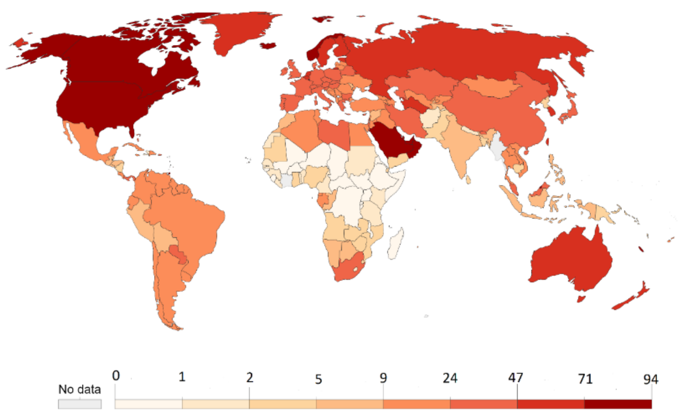

Source: adapted from Our World in Data[4]

The future – overcoming the limits of human brain power

Can we use the same approach that we used to measure progress and sustainability impacts of the industrial revolution (i.e., the concept of "energy servants") for understanding AI? In other words, how many AI servants – in human equivalents – are already amongst us?

Narrow AI that is specialized on performing specific tasks is similar to the energy converters providing artificial muscle power. The latter are also all specialized and there is no machine that could perform all of human's physical labour – yet together they amount for enormous physical labour capabilities. However, we cannot use primary energy used for AI and related technologies to estimate the number of AI servants already in the midst of us. This is because the efficiency of the predominant AI today – namely deep neural networks with supervised learning – are on the order of 10,000 to 100,000 times less energy efficient than the human brain (more on this later!). The brain runs at 16W at rest which goes much higher for analytical thinking activities, such as chess (hence its characterization as sport). In comparison, training a single state-of the art DNN to do human level visual identification requires around 650MWh – much more than a human brain.

Hence, the AI space is far from a truly virtual space of information and data. Data movements are also fundamentally limited by energy. Modern-day data centres are testament to this, with some of them having their own power plants. In any case, the artificial intelligence revolution, modern energy access will be a prerequisite for commanding an increasing army of hidden online servants and brain power.

However, we can start by comparing computational performance and data flows. While again not a perfect measure, it can provide a big picture of how far down the world is already in deploying AI and what to expect in terms of wider impacts. By some estimates, 20 Petaflops might roughly be the hardware-equivalent of the human brain (Kurzweil, 1999; Diamandis&Kotler, 2015).[5] Please note that one Petaflop is 10^{15} floating point operations per second (FLOP). It is a measure of computer performance. For comparison the world's top supercomputer in December 2019 had a performance of 201 Petaflops, similar to 10 human brains. By Nov. 2020, the Japanese Fugaku supercomputer reached a new world record of 442 Petaflops, or the equivalent of 22 human brains.

In other words, in order to have 21 AI servants per person similar to the 21 energy servants providing artificial muscle power, we would need computing capacity similar to one Fugaku supercomputers for every person on the planet. By the way, the top 500 supercomputers combined reached the equivalent of 50 human brains by the end of 2019 and more than 100 by

[4] https://ourworldindata.org/energy-production-consumption

[5] Please note this estimate is highly uncertain. A literature survey showed a range from 10^{12} to 10^{28} FLOPS https://aiimpacts.org/brain-performance-in-flops/ When we look at typical AI tasks and compare it to DNN implementations, the estimate of 20 Petaflops appears reasonable, though.

Nov. 2020. At current rates, the top supercomputer is expected to reach 1.2 million Petaflops or the performance equivalent of 58,000 brains by 2030 - the SDGs timeline. This would be another factor 2,600 improvement over a decade!

No official statistics exist for the number of FLOPS by all the world's computers, smart phones and other devices – most of which are connected to the Internet. This collective computing power (or computing capacity) worldwide was estimated at around 850,000 ± 650,000 Petaflops around the year 2015.[6] This is the hardware equivalent of 43,000 ± 32,000 human brains. While still small compared to the world population, this performance has the potential to continue growing exponentially, as it did in the past. In fact, it is estimated that global computing power increased almost by a factor 10,000 from 2007 to 2015 – which was roughly equally due to widespread digital technology adoption and performance increases of the devices.

Most recently, the performance of supercomputers and mass-produced computers has increased by compound annual rates of 79% and 33%, respectively. Considering continued growth in digital devices[7] and movement towards big data centers and cloud computing, reasonable estimates for the global computing power today (March 2021) is 93 ± 71 million Petaflops or 4.7 ± 3.6 million AI servants. By 2030, we might reach an estimated 150,000 Zettaflops or the human equivalent of 7.7 billion AI servants – basically a doubling in human cognitive capacity. By 2040, it might be as much as the human equivalent of 25 trillion AI servants (*Table 2. Numbers of "AI servants" deployed, 2010 - 2040*) – well beyond in magnitude to what modern energy converters managed to achieve in 200 years of industrial revolution.[8]

Alternative AI scenarios

I suggest the scenario analysis and integrated assessment communities to use the above simple framework to develop four global AI scenarios and to quantify their sustainability implications, in order to put the global debate on AI on a more systematic basis (see Table 3 for an overview of key scenario characteristics):

- Scenario 1: Dynamics as usual scenario
- Scenario 2: Sustainable development
- Scenario 3: Winner-takes-all
- Scenario 4: Regionalisation & fragmentation

The "dynamics as usual scenario" assume a continuation as in the past with the result that a majority of the world population might be left behind.

The "sustainable development scenario" foresees a truly multilateral world which promotes AI as a global public good, focused on SDG and SD achievement, leaving no one behind.

In the "winner-takes-all scenario", one country (and one or two companies) "win" the AI race, take all the benefits and use their new power to subdue or shut out everyone else (economically laissez faire).

The "regionalisation & fragmentation scenario" foresees a fragmentation of AI infrastructure, regulation and use into a few regional blocks, with little cooperation between them.

Table 2. Numbers of "AI servants" deployed, 2010 - 2040

	2010	2015	2020	2025	2030	2035	2040
World population [billion]	7.0	7.4	7.8	8.2	8.5	8.9	9.2
"AI servants" (human brain equivalents)	1	43,000	2 million	108 million	7.7 billion	420 billion	25 trillion
"AI servants" per person	~0	0.000006	0.0003	0.013	0.9	47	2700

[6] https://aiimpacts.org/global-computing-capacity/

[7] This assumes only a 10% compound annual growth rate in the number of digital computation devices, plus the most recent trends in performance increases.

[8] Note that these estimates are rather conservative, assuming a slowing increase in the number of devices – 8% per year until 2030 by which time most of the world population could be connected, and 3% per year thereafter in the 2030s, as the Internet of Things continues to expand.

Table 3. Overview of alternative AI scenarios

Scenario		Dynamics as usual	Sustainable development	Winner-takes-all	Regionalisation & fragmentation
Story		Continue as in the past, leaving most behind	A multilateral world promotes AI as a global public good, focused on SDG and SD achievement, leaving no one behind.	One country wins (and one or two companies) "win" the AI race, take all the benefits and use their new power to subdue or shut out everyone else (economically laissez faire)	Fragmentation of AI infrastructure, regulation and use into a few regional blocks, with little cooperation between them.
Values	Efficiency	High	Medium	High	Relatively high
	Equity	Low	Very High	Extremely low	Medium
	Autonomy	Very low	Medium	Extremely low	Medium
Governments	Int'l cooperation	Very low	Very high	Little to none	Globally very low, regionally high
	Role (vs. private sector)	Low, reactive (except in a handful of technological leaders), marginal in developing countries	High	Very high	Mixed bag, lots of PPPs
Technology	Rate of progress	Steady, medium, exponential, below potential due to concentration and little technology transfer.	Medium	Very high (due to scale economies)	Depends on region, low to high
	SD mission-oriented (geared towards more than efficiency)?	No, defence- and power-play driven	Very High	Irrelevant	only half-hearted, primarily in Europe
Sharing	Leaving no one behind	Not an agreed objective	Overarching objective	Irrelevant (the opposite power grab is the objective)	only half-hearted, primarily in Europe
	Ethical guard rails	Proliferation of frameworks with little real impact	Universally accepted and policed	Irrelevant	Europe and some major MNEs

Conclusion

In conclusion, important lessons can be learnt from the industrial revolution – not only in terms of the role of energy and technologies, but also in terms of social impacts and ethical considerations of balancing the objectives of efficiency with equality and autonomy. In particular, much can be understood from taking a broad picture view of AI, in which "AI servants" - each a hardware performance equivalent of a human brain – become available and partly complement and partly replace human cognitive tasks, greatly increasing efficiencies. At present, the world commands an estimated 4.7 million such AI servants - expected to rise 7.7 billion by 2030 and an incredible 25 trillion by 2040. This is in analogy to the industrial revolution, where human muscle power was increasingly replaced and complemented by energy converting machines doing the physical labour for us – in specific, well-defined areas. Today, everyone of us commands an average of 21 such "energy servants" doing physical labour for us.

One result is that the current AI transformation is about seven times as fast as the industrial revolution which has important lessons for policy response and the design of institutions.

Socio-economic divides arise from unequal ownership over the "AI servants and "energy servants".

Navigating a future for beneficial AI will require balancing the actions in support of common values of freedom/autonomy, human well-being/quality of life, and justice/equality.