Photo Credit: Von Wong https://www.vonwong.com

This Think Piece provides a meta-analysis of the environmental challenges presented by emerging technologies collectively referred to as *machine learning* (ML) and *artificial intelligence* (AI). The Think Piece is a call for action for the SDOs, Governments, Policy Makers & Regulators, System Developers & Business Organizations, Academia & Researchers, and Intergovernmental & Non-Governmental Organizations. In this think piece, we argue that AI poses collective challenges like the widespread introduction of plastics half a century ago - providing cheap and/or seemingly efficient solutions to a wide range of contemporary problems while creating (and displacing) new aggregate costs that will impact all of society. We argue that a holistic ethical and policy approach needs to be taken towards the design and deployment of AI technologies. It is imperative to embed an awareness of key ethical and policy goals throughout the ecosystem and value chain - from researchers to engineers to corporations to global policymakers - to better generate and distribute metrics that can inform policies oriented towards ethical and sustainable practices.

The emergence of disruptive digital technologies has already begun profoundly reshaping our lives, our interactions, and our lived environments. As adoption of these technologies becomes widespread, they are likely to play a substantial independent role in society's energy consumption and environmental impact.

Following the conclusion of the COP26 climate conference, private organizations and governments alike are stepping up their promises to combat climate change, bringing to bear a mix of public policy and innovative technologies to address one of our era's defining challenges.

It is well understood and acknowledged that Artificial Intelligence leverages predictive models to help scientists understand how climatic shifts might manifest across the world in the coming decades. But as exciting as it may be to contemplate a world where AI helps tackle the climate crisis, there's no escaping the bitter irony that AI itself comes with a significant carbon footprint.

Generally, the use of AI is synonymous with good innovation and positive applications of this technology. AI itself will also become smarter with each passing year. As we look forward to the gains, efficiencies, and new solutions that AI creates for nations, businesses, and everyday life, we must also think about how to maximize gains for society and our environment. The fact that AI will have a major impact on society is no longer in question. Current debate turns instead on how far this impact will be positive or negative, for whom, in which ways, in which places, and on what timescale.

Sustainability as an ETHICAL concern in AI

Sustainability calls for development and deployment of AI to consider protecting the environment, improving the planet's ecosystem and biodiversity, contributing to fairer and more equal societies, and promoting peace. Ideally, AI should create sustainable systems that process data sustainably and whose insights remain valid over time. To achieve this aim, AI should be designed, deployed, and managed with care to increase its energy efficiency and minimize its adverse ecological footprint. To make future developments sustainable, corporations are asked to create policies ensuring accountability in the domain of potential job losses and to use challenges as an opportunity for innovation.

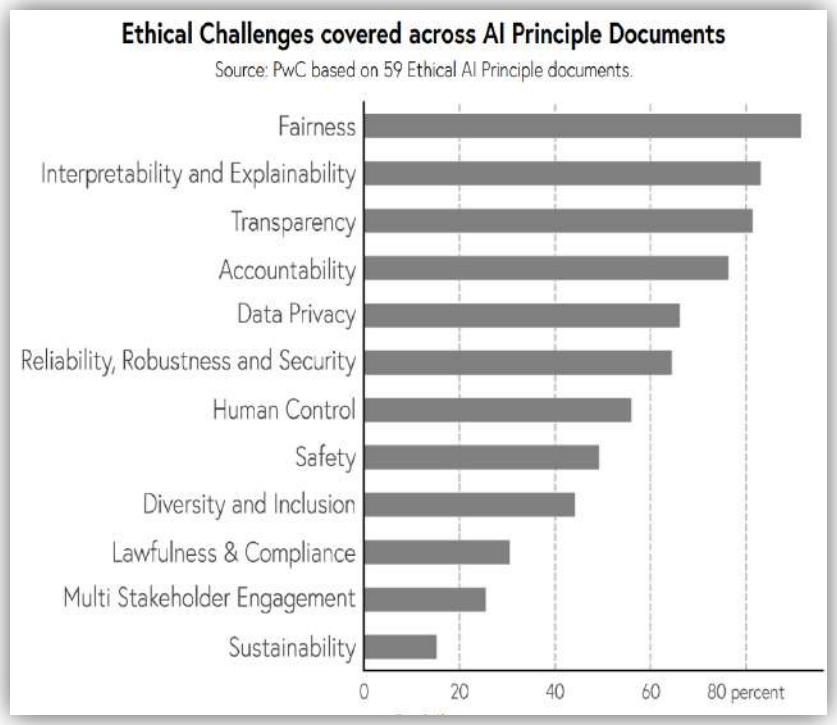

The relative thematic underrepresentation of sustainability and solidarity suggests that these topics might be currently flying under the radar of the mainstream ethical discourse on AI. While the attempts to justify it could be made on the pretext that the people who work on AI are usually not climate scientists The under-representation of sustainability-related principles is particularly problematic because the deployment of AI requires massive computational resources which, in turn, require high energy consumption. The environmental impact of AI, however, does not only involve the negative effects of high carbon-footprint digital infrastructures, but also the possibility of harnessing AI for the benefit of ecosystems and the entire biosphere.

A survey and analysis conducted by Jobin et al on the global landscape of ethics guidelines in AI found that only 14 documents out of 84 consider sustainability, environment (nature), energy, resources (energy) to some extent.

THE GLOBAL CHALLENGE: AI AND PLASTIC

How does AI work like plastic? It promises to be a cheaper and more effective alternative to business procedures that are commonly in place. But it has the potential to displace explicitly evident costs (fewer staff, more productive machinery, and procedures) into costs that are harder to perceive (massive energy costs generated at computational facilities). Also, will its implementation change the structure of our economy in ways that facilitate downstream burdens on the environment that are currently hindered by the limitations of our economy?

We need to learn from our past mistakes, else the adage - *history repeats itself* shall be proven true in next few decades (if *not* in years). Past failures to consider environmental consequences early have been costly: semiconductor industry (metals, solvents); synthetic chemicals (PCB, DDT, Freon); Applications of natural compounds (chlorine, asbestos); transportation, energy (air pollution, global warming, nuclear wastes).

THE ECONOMIC CHALLENGE: AI AND THE ENVIRONMENT

The artificial-intelligence industry is often compared to the oil industry: once mined and refined, data, like oil, can be a highly lucrative commodity. Now it seems the metaphor may extend even further. Like its fossil-fuel counterpart, the process of deep learning has an outsize environmental impact.

Although AI presents transformative opportunities to address the earth's environmental challenges, left unguided, it also has the capability to accelerate the environment's degradation.

There is a major hype to leverage AI to bring sustainability in different aspects forgetting the extensive energy it consumes, in turn increasing GHG emissions and *carbon footprint* it adds to anything, when you apply AI to any product, system or solution.

Recent analyses of NLP training processes suggests that the energy required to train an NLP model may be substantially higher than widely recognized. Do we have mechanisms to compare the gross savings of efficiency that AI systems will provide against the gross cost of their training relative to their anticipated life span (to provide a net efficiency analysis)? Also, cost of inference (deployment computation) vs cost of training (development computation)?

THE POLICY CHALLENGE

To address the ethical implications - stakeholders must be identified, including the relative allocation of threats/harms and promises/benefits among them.

Local Magnitudes: Do system-level metrics adequately account for uneven distributions in magnitude of harms/benefits. Does distribution of harms unduly burden vulnerable communities? Do benefits exacerbate controversial social inequalities?

Impersonal Interests: Are there abstract or non-human interests (eco-systems, animals, etc.) that deserve independent ethical consideration? If so, how should they be balanced against the needs of human communities?

Identify Stakeholders: Who bears the environmental threat/harm of computational carbon cost, and are those threats/harms captured by metrics? Who stands to benefit?

RESOURCE & ENERGY INTENSITY OF ARTIFICIAL INTELLIGENCE

In recent years, AI has increasingly meant using deep learning models that require large amounts of computing and electricity to run, such as OpenAI's GPT3 language model, trained on almost 500 billion words and using 175 billion parameters.

The discussion about AI opportunities and risks has only recently begun to consider how much energy and resources AI itself consumes for computing. The training period of artificial neural networks (ANN) devours particularly large amounts of energy. A study using BERT, an ANN used for speech recognition, found that the training period alone resulted in 0.65 tons of CO_2 being emitted. This amount

corresponds to the emissions generated from a return flight between Berlin and Madrid…

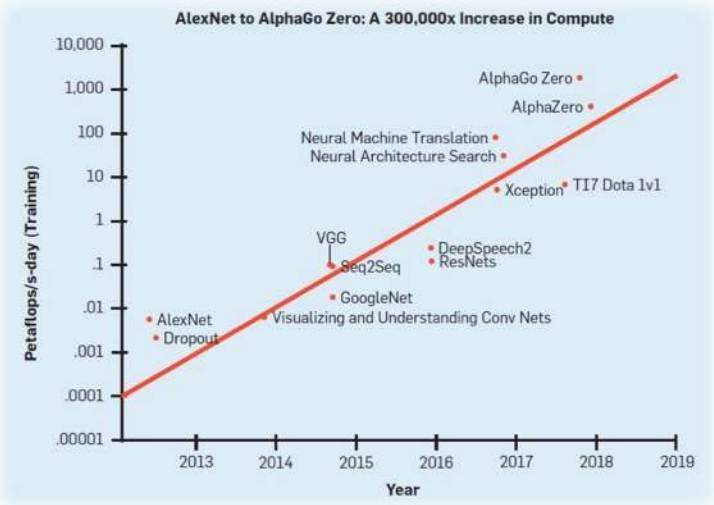

The amount of compute used to train deep learning models has increased 300,000x in six years.

RED AI v/s GREEN AI

Artificial intelligence promises both to combat the ill effects of climate change and to make the emergency worse with excessive energy use. How can AI be made greener?

Implementation of an AI solution is often preceded by multiple iterations of refinement by tuning hyperparameters. Needless to say, this can incur a large computational cost, and consequently, a large environmental cost. Creating efficiency in AI research will decrease its carbon footprint and increase its inclusivity as deep learning study should not require the deepest pockets.

The term Green AI refers to AI research that yields novel results while considering the computational cost, encouraging a reduction in resources spent. Whereas Red AI (AI development that is oblivious to its environmental impact) has resulted in rapidly escalating computational (and thus carbon) costs, Green AI promotes approaches that have favourable performance/efficiency trade-offs.

The efficiency of Red AI solutions can be improved as follow-up work, and this can be achieved by training on a smaller dataset. The computational cost of recent work, they feel, pays off in downstream performance. However, this rate of return is declining with time.

A Green AI approach is where both training and inference aspects of running machine learning algorithms are optimized for computational efficiency. This is necessary because models used in productions incur higher computational costs during inference, whereas in research training is done much more frequently.

It is imperative to benchmark Efficiency across four categories: *Training Time; Training Cost; Inference Latency; Inference Cost.*

WHAT IT WILL TAKE TO REDUCE AI'S CARBON FOOTPRINT

Shrinking AI's carbon footprint is only possible if we first understand the scope of the problem. Fortunately, there are steps tech industry leaders can take to ensure that AI innovation doesn't come at the expense of the planet's health. From rethinking hardware and the complexity of models to reducing processing required in both the

training and inference stages, here's what it will take to achieve eco-friendly AI innovation.

NO to Power-Hungry Models; Work Smarter, NOT Harder; Smaller is Greener

METRICS TO MEASURE THE ENVIRONMENTAL COSTS OF AI COMPUTATION

The precise environmental cost of AI solutions is not easily determined. AI processes may displace known costs in unexpected ways. Environmental costs create a distribution of harms and benefits that have ethical implications. It has been found historically that the impact of deteriorating environmental conditions is disproportionate. The brunt of the environmental costs is borne by low income, minority communities, who are already marginalized. How those costs are measured can render some harms invisible or create misleading senses of scale for others. Measurement itself has an ethical dimension. Aligning AI to human values means picking the right metrics

Types of Costs associated with AI

Run Cost; Production Cost; Expansion Cost; Financial Cost as a proxy for environmental consumption; Electric Consumption, Net Operations/Heat Production.

The ability to measure and track AI's carbon footprint is a step in the right direction, but the industry on a whole, needs to be doing more. Thankfully, there are steps that can be readily implemented.

The good news is that there are ways to optimize the core architectures of deep learning models that can increase performance efficiency without detracting from their accuracy.

The aim of an eco-label for resource- and energy-efficient software products is to reduce the overall energy consumption of information and communication technology and to increase resource efficiency.

However, it would be important to note that the current development in AI is in many cases characterized by aggressive global competition for dominance in AI due to the "winner takes it all" situation on the market.

Measuring and Reducing AI's Carbon Footprint

There is a real need to think about how organizations are building the AI systems. Are they training a needlessly complex algorithm? How frequently are they retraining? There's also the AI supply chain, procuring algorithms, procuring computational hardware, and thinking about the carbon footprint.

AI has both a negative and positive effect on the environment, and it's important to measure both. Many organizations feel the pressure of The Paris Agreement that requires member states to reduce their carbon footprint by 55% by 2030, compared to 1990. This pressure has been re-affirmed in the COP 26 by all nations and organizations by providing the timelines to achieve Net Zero Carbon Footprint.

If you're going to run AI, you're going to need 'the machines,' so you're going to leave a carbon footprint. As you get more data and the models [become] more complicated, the more energy you're going to consume to get the AI models you're looking for.

But how can a data scientist measure AI's footprint?

Tools and frameworks are emerging to help understand AI's impact. Meanwhile, technology continues to improve at every layer of the tech stack, which should also help reduce carbon emissions, from algorithms to data centres and beyond.

Develop Enabling Technical Environments for The Green Transition

Technologists need to apply their skills to climate change mitigation and work towards transforming how data-driven solutions are being developed and commercialized at scale. Industries like energy, food, manufacturing, and finance need to transition within the next five years.

Develop a Climate Aware Data Science Practice

AI and data science communities will need to follow the steps of computer scientists who have a long history of investigating sustainable computing. Researchers may advocate for making

efficiency an evaluation criterion for research, use more computationally efficient hardware and algorithms and report the "price tag" of their models. Alternatively, Energy Usage Reports have been proposed as part of greener algorithmic accountability practices and tools like Machine Learning Emissions Calculator can help estimate the amount of carbon emissions produced by the training of AI models. Similarly, practitioners may start reporting the time to retrain models, share local infrastructure instead of relying on cloud computing and choose cloud providers who are offsetting their emissions.

Focus on Climate Justice

A just transition requires that we pay attention to the struggles of various communities who are already defending their land, air, water, and livelihoods from extractive activities such as mining, fracking, gas flaring, etc. Any application of AI in climate change mitigation and adaptation will need to ensure that environmental impacts are not externalized onto the most marginalized populations, and that the gains are not only captured by digitally mature countries in the global north. This requires centering front-line communities and enabling them to take ownership of their data and bottom-up climate action plans.

STANDARDS: AN ORGANIZATIONAL SOLUTION

And this is where STANDARDS can really help. There is plethora of standards to help us develop most Optimized, Energy Efficient and Environment friendly products, systems, solutions and even services. We also have Sustainability standards, Frameworks and Metrices to measure the extent of sustainability, circularity and carbon footprint of any product, service and/or organization itself.

Standards: There is no common standard for evaluating the environmental cost of AI systems, but a diverse range of methods are being developed.

RECOMMENDATIONS: CALL FOR ACTION

The Sustainability in context of the Artificial Intelligence is truly a crosscutting subject for all the stakeholders, be it the SDOs, the Governments, the Policy Makers & Regulators, the System Developers & Business Organizations, the Academia & Researchers, or the Intergovernmental & NGOs. To ensure that AI does not lead to environmental disaster like we are witnessing because of thoughtless proliferation of Plastic, it would be imperative to work in a close co-ordination and collaboration by all the stakeholders.

AI technology also has the potential to amplify and exacerbate many of the risks we face today. To be sure that AI is developed and governed wisely, government and industry leaders must ensure the **Sustainability** ***(energy efficiency, circularity, and carbon footprint)*** along with the already accepted aspects like Safety, Explainability, Transparency and validity of AI applications. It is incumbent on authorities, AI researchers, technology pioneers and AI adopters in industry alike to encourage deployments that earn trust and avoid abuse of the social contract.

Achieving this requires a collaborative effort to ensure that as AI progresses, its idea of a good future is aligned to human values and encapsulates a future that is safe for humanity in all respects – its people and their planet. It would be imperative to consider Sustainability as a moral framework based on social justice, which can be used to evaluate technological choices.

Though, for last one decade, we have been witnessing isolated initiatives in different domains like Green ICT, Inexact Computing, Green Data Centres etc. it is time to consolidate such siloed efforts into a comprehensive crusade to understand and mitigate the looming calamity that disruptive technologies bring along with systems thinking and systemic approach. The fragmented landscape of sustainability standards should be modernized to get the most benefit from the digital revolution to help achieve climate and sustainability goals.

CONCLUSION

The combined effect of climate change and society's impact on the earth is intensifying struggles over natural resources while also threatening our infrastructure, food systems and quality of life.

It has been observed that the technologies developed by human beings in the last 2-3 centuries have had a major impact on the earth's climate and our nature's equilibrium. Some believe that we have reached a point of no return. This can have a huge impact on life on earth, especially on the human species.

However, while technology has been responsible for most of it, technology also seems to have a solution for it. It's increasingly clear that today's environmental conditions are not sustainable. Over the past few decades, we have made huge progress fighting disease, poverty, and illiteracy. Now we must apply that same ingenuity to the problem of global warming and other consequences of human activity.

AI systems, and their ability to control machines automatically and remotely, have caught the public's imagination. The opportunity for AI to be harnessed to benefit humankind and its environment is substantial. AI as a technology can help to mitigate climate change. This aspect should be further emphasized in terms of measuring the benefits of AI for the environment and climate. The intelligence and productivity gain that AI will deliver can unlock new solutions to society's most pressing environmental challenges: climate change, biodiversity, ocean health, water management, air pollution, and resilience, among others.

In a world increasingly dominated by AI applications, an understudied aspect is the carbon and social footprint of these power-hungry algorithms that require copious computation and a trove of data for training and prediction. While profitable in the short-term, these practices are unsustainable and socially extractive from both a data-use and energy-use perspective.

Sustainability must be at the centre of a digital strategy – not a sideshow or an afterthought. If sustainability is part of the digital business case and reason for investing, it will be part of how targets are set, and value is measured.

Give technology new roles to play - The next step is to reimagine the role of digital technology. Until now, it's largely helped to boost efficiency, cut costs, and promote consumption. Now it must focus on comprehensive sustainability.

We humans need to radically change our relationship, not just with the planet, but with the objects with which we fill our lives. We need to change how we think about technology and innovation. Rather than allowing technological advancement to steer our narratives, innovation and technology should help us build bridges between the worlds we inhabit now and the ones we imagine for tomorrow.

This "Think Piece" is developed by OCEANIS (Open Community for Ethics in Autonomous and Intelligent systems) - A Global Forum for discussion, debate and collaboration for organizations interested in the development and use of standards to further the development of autonomous and intelligent systems. Author (N. Kishor Narang) is the Chair of 'Advancing Research Work Group' in OCEANIS and leading the initiative of identifying gaps, establishing policies, and adapting new technology solutions that will solidify global economic prosperity by driving important research and policy to help close some of the most significant gaps in technology expansion.