Position: Who Gets to Harness (X)AI? For Billion-Dollar Organizations Only

Jonathan Dodge[a]

[a]Oregon State University, 1148 Kelley Engineering Center, Corvallis, OR 97331-5501

Abstract

Recently, researchers have made a number of tremendous advancements in AI capabilities. However, I will show that the fixed financial cost of making such advancements is high, and further, so are the recurring energy costs. As a result we see "AI haves and have nots" with wildly differing amounts of power between these two groups. This means our research community needs to carefully consider the role XAI has in mediating communication between stakeholders in the face of such an important power dynamic. This paper aims to engage that process, examining the current state of affairs through a variety of lenses, and then identifying some promising ideas for the future.

Keywords

Explainable AI, Social Aspects of AI, Social Aspects of Explanation

Introduction

This paper considers each of the following in turn: (1) What does AI cost to setup/scale? (2) What does energy to run hardware for AI cost? (3) As a result of 1 and 2, who is getting left behind? (4) What's the alternative?

Explanation's role largely appears in the third question, where I argue that as a downstream consequence of inequity of access to AI, explainable AI will suffer a similar fate.

1. What does AI cost to setup?

A lot. According to the 2020 State of AI Report [1], Microsoft-backed OpenAI likely spent more than **$10M** training GPT-3, a state-of-the-art natural language processing system. In justifying the decision to release as a commercialized API rather than open-source, OpenAI states (underlining added for emphasis):

> *…many of the models underlying the API are very large, taking a lot of expertise to develop and deploy and making them very expensive to run. <u>This makes it hard for anyone except larger companies to benefit from the underlying technology</u>. We're hopeful that the API will make powerful AI systems more accessible to smaller businesses and organizations."* [2]

To cite another example, while AlphaZero only requires 30 hours [3] of training (using massive parallelism) to achieve a skill level exceeding the version of AlphaGo that defeated Go master Lee Sedol[1], the journey to the destination was not so cheap. Wired reported an estimate[2] that duplicating the training time of the experiments leading to that

Joint Proceedings of the ACM IUI 2021 Workshops
dodgej@oregonstate.edu (J. Dodge)

[1]Lee Sedol would go on to retire, stating *"Even if I become the number one, there is an entity that cannot be defeated."* [4].

[2]While the underlying source is not necessarily reputable—a blog post [5]—they present the full chain of reasoning to arrive at this cost.

model would cost **$35M** [6]. Medium [7] arrives at a lower estimated cost to train the more complex AlphaStar—**$12M**—but that estimate is just to replicate the final agent(s), and does not include formative experiments. Of note, part of these cost estimates include renting time on proprietary Tensor Processing Unit (TPU) hardware [8], so the internal costs to DeepMind would be lower than the estimates presented here. Of additional note, the Wired article [6] opens with scary points about how much money DeepMind has lost (more than **$1B** in 3 years), and develops an argument about how these steep losses may chill investment in AI research.

One might expect high costs to yield effective solutions, if not viable products. However, in the face recognition market, IBM, Microsoft, and Amazon have retreated (at least temporarily), potentially abandoning their investments. To understand why, here are excerpts from each of their public statements:

> *"IBM no longer offers general purpose IBM facial recognition or analysis software. IBM firmly opposes and will not condone uses of any technology, including facial recognition technology offered by other vendors, for mass surveillance, racial profiling, violations of basic human rights and freedoms, or any purpose which is not consistent with our values and Principles of Trust and Transparency."*
>
> —Dr. Arvind Krishna, IBM CEO, June 8, 2020 [9]

> *"We're implementing a one-year moratorium on police use of Amazon's facial recognition technology. ... We've advocated that governments should put in place stronger regulations to govern the ethical use of facial recognition technology, and in recent days, Congress appears ready to take on this challenge. We hope this one-year moratorium might give Congress enough time to implement appropriate rules, and we stand ready to help if requested."*
>
> —Amazon Staff, June 10, 2020 [10]

> *"But I do think this is a moment in time that really calls on us to listen more, to learn more, and most importantly, to do more. Given that, we've decided that we will not sell facial recognition to police departments in the United States until we have a national law in place grounded in human rights that will govern this technology."*
>
> —Brad Smith, Microsoft President, June 11, 2020 [11]

These rich and powerful companies seem to feel facial recognition is presently too legally and/or ethically thorny for at least some applications, e.g. law enforcement. Or possibly it is too difficult to do with sufficient robustness using current techniques and data, as indicated by findings like the following, from Buolamwini et al., *"All classifiers perform worst on darker female faces (20.8%-34.7% error rate)"* [12]. Either way, how can one imagine similarly large problems as tractable for organizations with smaller budgets, development teams, and legal departments?

So what if one wants to do something on a smaller scale, by purchasing a GPU? Tim Dettmers, currently a PhD student at University of Washington, has hosted a benchmark comparison of various cards since 2014. The September 2020 revision[3] recommends RTX 3080 or RTX 3090 [14], which has a retail price of around $700. In his post's TL;DR advice, he offers the following:

> ***I have little money:*** *Buy used cards. Hierarchy: RTX 2070 ($400), RTX 2060*

[3] Notably, this revision was made in the middle of the pandemic, and prices have gotten *much* worse since, if one can even *find* a card to purchase [13].

($300), GTX 1070 ($220), GTX 1070 Ti ($230), GTX 1650 Super ($190), GTX 980 Ti (6GB $150).

While this may not sound like a lot of money to some readers, as we will see in Section 3, it *is* for many people. Unfortunately, this does not even include the cost of the *rest* of the computer (e.g. motherboard, CPU, peripherals for I/O). Nor does it account for energy to run the system.

2. What does energy to run hardware for AI cost?

A lot. According to NVIDIA's tech specs, the RTX 3080 mentioned earlier consumes **320W** *by itself*, thus requiring at least a **750W** power supply for the full system [15]. A full breakdown of recommended power supplies for CPU plus GPU pairings comes courtesy of ASUS, shown in Figure 1. Note that the upper diagonal has power draws that are *twice* those in the lower diagonal, and that the newest generation of GPUs use much more power than previous generations.

So now armed with estimates of power consumption required by some relevant hardware, next we investigate length of training times:

> *"[After starting with 64x64 images] ...we have ramped up to bigger images, it takes a while, it takes* ***2-3 days*** *to train a GAN. We don't do grid search, but even random search you still need to train 12 models times however many configurations you have, so it really adds up."*
>
> —Dr. Sasha Luccioni, Postdoc in AI for Humanity [17]

Notably, the GAN Dr. Luccioni refers to is intended to visualize climate change effects, so training efficiency is of particular concern to her. Further, generating one production model requires a great deal of prior experimentation. To illustrate, Hill et al. report:

> *"All 7 Interview #2 respondents reported evaluating the model to be a long, arduous process, in which participants iteratively refined the model through changes to parameters, training data, and so on. They did these iterations using an evaluate-fix-evaluate cycle that often required many changes, sometimes even sending the participants back to earlier steps..."* [18]

To illustrate the length of training times needed for certain challenge problems, during AlphaStar's demonstration matches against professional StarCraft II players, one of its creators made the following remarks:

> *"The league here was run for about* ***a week****... 7 days of real time is actually longer in StarCraft, as Tim was saying we got a binary that can run the game much faster. So the most experienced agents we see today have played about* ***200 years*** *of StarCraft II."*
>
> —Dr. Oriol Vinyals, Research Scientist at Deepmind [19]

As a sidenote, the league used to train the agent pool contains a few hundred agents, with Vinyals et al. [20] stating *"During league training almost 900 distinct players were created"*. Consider the Cartesian product involved with these numbers: 900 agents × 200 years × 525600 minutes/year × 280 actions/minute ≈ 26.5 **trillion** actions selected[4]. To produce the huge quantities of floating point operations to do this (and the necessary experimentation to devise the final training process), it is important to realize that processes like

[4] Action per minute (APM) estimate is based on data presented in [19]. The rest of the numbers used in this estimate are natural constants or from quotes earlier in the paper.

RECOMMENDED PSU TABLE

Please refer to the recommended PSU table as below. It indicates clearly on the power by GPUs and corresponding CPUs. We hope this information help you choose the right power supply units based on your components.

NVIDIA GPU - Ampere Series

	Intel I5 AMD Ryzen5	Intel I7 AMD Ryzen7	Intel I9 AMD Ryzen9	Intel HEDT AMD ThreadRipper
RTX 3090	750W	750W	850W	1000W
RTX 3080	750W	750W	850W	850W
RTX 3090	650W	650W	750W	850W
RTX 3090	550W	650W	750W	750W

Turing Series

	Intel I5 AMD Ryzen5	Intel I7 AMD Ryzen7	Intel I9 AMD Ryzen9	Intel HEDT AMD ThreadRipper
RTX 2080 TI	650W	750W	750W	850W
RTX 2080 Super	650W	650W	750W	850W
RTX 2080	650W	650W	750W	850W
RTX 2070 Super	550W	650W	650W	750W
RTX 2070	550W	650W	650W	750W
RTX 2060 Super	550W	550W	650W	750W
RTX 2060	550W	550W	650W	750W
GTX 1660 TI	450W	450W	550W	650W
GTX 1660 Super	450W	450W	550W	650W
GTX 1660	450W	450W	550W	650W
GTX 1650 Super	450W	450W	450W	550W
GTX 1650	450W	450W	450W	550W

Figure 1: (Source: ASUS [16], redrawn for better printing) Recommended power supply wattage for different CPU+GPU combinations. Note how much power new generation components draw when compared with even recent generations of hardware by comparing down columns (basically adds **100W** across the board, up to **300W**).

these will involve running *many* large capacity servers for *long* periods of time.

Unfortunately, it is hard to know the energy costs of many AI systems because of obscured costs:

> *"It's really really hard to quantify exactly how much energy you are consuming when training a neural network because often it is on the cloud, or you are sharing a GPU with others..."*
>
> —Dr. Sasha Luccioni [17]

Google's proprietary TPU obscures these costs further because its power consumption is unknown as of the time of this writing, but estimated at **200-250W** by [21].

Training on the cloud has advantages beyond convenience, but first some context:

> *"Depending on the grid you are connected to, you know the energy mix of the grid... Like in the U.S. there's, I don't know, 20-30 grids, some are really clean, and some are really coal based... So depending where [your training] is, your carbon emissions can vary up to 80x."*
>
> —Dr. Sasha Luccioni [17]

So, in effect, if one's local grid runs on coal, training on the cloud might be a way to get long-running processes onto a cleaner grid, e.g. one running on hydro power. To do so, all that would be needed is to build data centers near power plants, particularly clean ones. Note that this is already under proposal in some localities (e.g. [22]).

However, training on the cloud also has

a major drawback—obscured costs. Previous research in resource consumption awareness by Strengers found that people struggle to understand resource management units, preferring visual analogs, e.g. buckets of water [23]. For water, the hard-to-understand flow rate (e.g. gallons/min) has an *easy*-to-understand volume analog (gallon)[5]. Contrastingly, both electricity's flow rate unit *and* "volume analog" are hard to understand (e.g. watt and watt hour, respectively). While most understand that kilowatt hours appear on power bills to measure payment, few can conjure a visual or describe the meaning of the term.

All of this is to say that electricity consumption habits and their consequences are already obscure due to the nature of electron flow, and adding another layer of obfuscation—the cloud—makes it even harder to observe: 1) *that* consumption is occurring, 2) *how much*, and 3) the *consequences* of consumption. So what results from these factors? Inefficient consumption is a common result, i.e.:

> *"It's a problem in deep learning nowadays, 'Your model is not doing well?' 'Train on more data. Add more layers.' People's reflexes are essentially bigger, more... People who have been around in the field for a long time... their first reflex is* ***NOT*** *get more data, it is: 'Have you looked at your model? Have you figured out what it is doing? Do you know what is going on in these layers? Are you sure you have the right learning rate?' ... more about the fundamental stuff, whereas nowadays its like 'Throw more data at it'... because we can."*
>
> —Dr. Sasha Luccioni [17]

[5]To illustrate the difficulty of actually perceiving flow rates, consider the following (thought) experiment: Turn on your faucet a little bit so that the water is going right down the sink. Then, note the visual appearance, open the faucet more, and compare. Fairly large flow rate changes will look essentially the same until the water stream gets much wider or starts to roil.

3. Who is getting left behind in this?

Most everyone. A pre-pandemic survey by Pew found:

> *"61% of Americans say there is too much economic inequality in the country today... For those who say reducing inequality should be a government priority, large majorities point to unfair access it affords the wealthy and limits it places on others."* [24]

Since that study, equality of access has not improved. In a July 2020 publication (survey dates March 28 and April 4) by Bartik et al., small-businesses *"reported having reduced their active employment by 39% since January."* [25]. While there is a depth of research on wealth/income inequality[6], these will suffice to illustrate the point: Many businesses (and individuals) are priced out of AI, so are therefore priced out of XAI. This point is crystallized by Andreas Madsen, who first published as an independent researcher [27] but is now a PhD student at Université de Montréal, Mila. When asked to describe why he recommended that researchers *avoid* researching independently, as he just had done:

> *"If you just want to develop machine learning, like in industry, and you don't want to do anything novel, you can do that independently... But if you want to do research, if you want to develop new methods, this is a really hard field to do it in, because you need huge computational resources that you are not gonna have."*
>
> —Andreas Madsen [28]

We have argued previously [29] that, *"explanations are not just for people to understand*

[6]To the reader interested in these topics, I recommend [26], or the papers to which they respond.

the ML system, they also provide a more effective ***interface*** *for the human in-the-loop...".* If one considers that the explanation originator is an AI-enabled entity, they are likely a billion dollar company. On the other side of the explanation interface, if the explanation consumer is a member of the general public they are likely to be much poorer. Thus, XAI techniques mediate the interaction between two entities with a very lopsided power dynamic.

As a result of this power dynamic, the rich entity possessing the AI-powered system has an inherent conflict of interest. Consider four stakeholders[7]: Rich Uncle Pennybags (the character in the Monopoly game), Jon the XAI scientist, Alice the small business owner, and Joe on the street. While these various people have varying roles in the creation and consumption of explanations, Rich Uncle Pennybags is the *only* one with substantial power, and his incentive is to help himself economically. Critically, this leads to incentives that may *conflict* with the goals of explaining. For example, instead of explaining transparently, Rich Uncle Pennybags might harness explanations to protect the brand or to increase sales by omitting particular kinds of information. The analogy I would draw is that between peers, *"Because I said so"* would not be considered an acceptable explanation—though it might be as the power dynamic grows more lopsided (e.g. a boss might say that to their employee, or a parent to their child).

To illustrate the potential lopsidedness of power dynamics in XAI as interface, consider two scenarios. The first is a simple thought experiment: Imagine a median income wage-earner (e.g. Joe on the street) gets denied for a bank loan by an AI-powered system created by Jon the XAI scientist, on behalf of Rich Uncle Pennybags. As the would-be loan recipient consumes the explanation that they may have the right to receive in some jurisdictions, they will be attempting to understand the arcane decisions of an organization with power over them. While doing so, the user may suspect the service provider is just hiding behind the AI and/or its explanation. This is a dark pattern, described as "obstruction" by Chromik et al. [32]. Concerns surrounding abuse of dark patterns (intentional or otherwise) seem particularly pertinent given how many of the big AI players have previously been sued (or are currently being sued) for various anti-competitive practices. The list has grown to include (chronologically): IBM [33], Microsoft [34], Google [35], and[8] Facebook [37].

The second illustration comes from a discussion between TWIML/AI host Sam Charrington and guest Dr. Michael I. Jordan, a distinguished professor at UC Berkeley. There are two important aspects of this conversion (emphasized with boldface):

> **Jordan**: *"If you use my data, then that's ok with me if you use it and I get value out of it, if you pay me in some sense for using my data. In particular, a* ***travel agent*** *is a person who makes travel plans for lots and lots of people, and they get better and better at it over the years. ... So a system that does that, partially anonymizes me but that builds up experience dealing with people like me, then I kinda know what I'm getting. ... The medical system, you treat me and it works, I want it to be available to you tomorrow, I don't want to protect that data."*
>
> **Charrington**: *We think about AI in the context of this digital divide and there will be communities that will be left be-*

[7]This exercise is similar to the approach found in Ehsan et al's work, [30, 31], where they break down explainability as a socio-technical phenomenon focusing on *who* is doing *what*, *when*, and *why*.

[8]As of the time of this writing, lawsuits against Apple and Amazon have not been announced, though they were also involved in a recent congressional hearing on antitrust [36].

hind because they don't have ready access to AI technology. But the travel agent example makes me think that in a lot of ways, its like a human divide, in a sense, that we need to worry that the knowledge of humans is going to be commoditized into these computational systems, like AIs ... It gets really scary if its medicine, for example, where the masses are being treated by the ***commoditized robot doctors*** *but only the select few can speak to an actual human doctor*." [38]

4. Whats the alternative?

So it seems the (X)AI community hurtles toward a future where *few* can *afford* industry scale, with most relegated to toy problems.

4.1. Possible answer 1: Futurism

In their book[9], *"Why Can't We All Just Get Along?: How Science Can Enable A More Cooperative Future"*, authors Drs. Christopher Fry and Henry Lieberman offer a possible future paradigm **"Makerism"** in which everyone *"owns their own means of production"* [39]. They envision the advent of 3-D printing technologies, combined with advances in programming, energy generation, and materials recycling as enabling such a future.

A similar argument appeared in a recent Wired article entitled *"Want a More Equitable Future? Empower Citizen Developers"*. Authored by Microsoft CEO Satya Nadella and Harvard Professor Dr. Marco Iansiti, the article argues, *"Tools from cloud computing to AI should be in the hands of every knowledge worker, first-line worker, organization, and public sector agency around the world."* [40]. However, these optimistic visions do not seem to sufficiently address or resolve the existing economic imbalances described earlier.

4.2. Possible answer 2: Emerging Hardware/Software Architectures

While XAI is usually considered a software problem:

> *"It is important to understand that the reason why ML applications have received so much attention and widespread usage is because of the hardware."*
>
> —Dr. Diana Marculsecu, Chair of the ECE Department at UT Austin [41]

This paper already touched on the TPU architecture briefly, and it seems like they might be more power efficient than GPUs. But what other hardware might be possible? Dr. Marculsecu argues for a co-design process between software and hardware, which requires:

> *"...the ability to characterize neural network architectures, or perhaps components thereof, in terms of the time it takes to process those things, the energy or power it takes, and maybe come up with a joint metric that characterizes their energy efficiency or latency, per correct inference."*
>
> —Dr. Diana Marculsecu [41]

On the software side alone, a number of researchers are working to reduce power consumption. One reason this is necessary is that sometimes data must be processed on mobile devices which impose stiff constraints on computing and battery resources, e.g. for privacy and security reasons[10]. To take a single example, Bejnordi et al. propose a "channel gating" technique, in which the network

[9] I recommend this book highly, and hope the reader shares my appreciation for its game theoretic framing, via both prisoners dilemma and ultimatum game.

[10] Unfortunately, while each mobile device or IoT widget might consume very little power *individually*, these devices are so prevalent that *collective* energy costs will accumulate quickly due to the law of large numbers.

should learn when and which parts of the network can be short circuited [42]. Thus, researchers can train large capacity networks more efficiently, because, *"...there is no reason to compute features that help differentiate between several dog breeds, if there is no dog to be seen in the image."* [42].

Setting aside novel techniques, there are a wide variety of classical techniques that are generally considered to be more frugal than deep learning (e.g. random forests or Bayesian networks). However, since deep learning techniques currently dominate performance leaderboards, some of the cheaper techniques have fallen out of favor. Perhaps given hardware that can execute these architectures more efficiently they could be competitive?

4.3. Possible answer 3: Consumption Awareness and Disclosure

To raise awareness, many researchers are trying to make it easier for AI practitioners to estimate costs, both in terms of money and carbon footprint. Recent work by Strubell et al. [43] provides estimates of costs to train various benchmark models (see their Table 3). However, what if one has a custom model that differs from the benchmark designs by a great deal? Recent work by Cai et al. [44] aims to build predictive models capable of returning *"Detailed power, runtime & energy, with breakdowns"* given a convolutional neural network (CNN). Similar in spirit, Lacoste et al. recently proposed the "Machine Learning Emissions Calculator" [45], which provides a nice set of action items for individuals[11].

Last, when discussing energy consumption, one cannot overlook *"sample efficiency"*. Kevin Vu [47]'s Table 2 explains this by illustrating that lower wall clock time to solve does not always coincide with the fewest iterations, because various techniques require different amounts of time *per iteration*. The spirit of this branch of research is well encapsulated by a discussion section question of Henderson et al., *"How much is your performance gain worth? Balancing gains with cost"* [48].

4.4. Possible answer 4: Ethics training and practice among AI creators:

Failing any advances in equitability, improving the foresight of AI practitioners is a good second best.

> *"We are creating all kinds of artifacts, including social networks, that look fun... but it has turned out to not be so great in various ways... So a field that builds those systems <u>should also be aware of the consequences</u> of building that."*
>
> —Dr. Michael I. Jordan [38]

Possibly to similar ends, ACM revised its code of ethics in 2018 [49]. The first bullet point in the code is, *"Contribute to society and to human well-being, acknowledging that <u>all people</u> are stakeholders in computing."* Unfortunately, not all practitioners fully engage with these concepts:

> *"How do you know the unknowns that you're being unfair towards? [...] <u>You just have to put your model out there, and then you'll know if there's fairness issues if someone raises hell online</u>."*
>
> —Participant R7, Software Engineer [50]

4.5. Possible answers 5+: Choose your own adventure!

Plenty of ideas went unexplored here, which do you find promising, dear reader?

[11] If the reader wishes to engage in some AI research in the climate change space, I recommend [46], as they attempt to identify a variety of research gaps.

References

[1] N. Benaich, I. Hogarth, State of ai report, 2020. URL: https://www.stateof.ai/.

[2] OpenAI, Openai api faq, 2020. URL: https://openai.com/blog/openai-api/.

[3] D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel, T. Lillicrap, K. Simonyan, D. Hassabis, A general reinforcement learning algorithm that masters chess, shogi, and go through self-play, Science 362 (2018) 1140–1144. URL: https://science.sciencemag.org/content/362/6419/1140. doi:10.1126/science.aar6404.

[4] CNN (James Griffiths), Korean go master quits the game because ai 'cannot be defeated', CNN, 2019. URL: https://www.cnn.com/2019/11/27/tech/go-master-google-intl-hnk-scli/index.html.

[5] D. H, How much did alphago zero cost?, 2020. URL: https://www.yuzeh.com/data/agz-cost.html.

[6] Wired (Gary Marcus), Deepmind's losses and the future of artificial intelligence, 2019. URL: https://www.wired.com/story/deepminds-losses-future-artificial-intelligence/.

[7] Medium (Ken Wang), Deepmind achieved starcraft ii grandmaster level, but at what cost?, 2020. URL: https://medium.com/swlh/deepmind-achieved-starcraft-ii-grandmaster-level-but-at-what-cost-32891dd990e4.

[8] Google, Cloud tpu, 2020. URL: https://cloud.google.com/tpu.

[9] A. Krishna, Ibm ceo's letter to congress on racial justice reform, 2020. URL: https://www.ibm.com/blogs/policy/facial-recognition-sunset-racial-justice-reforms/.

[10] Amazon Staff, We are implementing a one-year moratorium on police use of rekognition, 2020. URL: https://www.aboutamazon.com/news/policy-news-views/we-are-implementing-a-one-year-moratorium-on-police-use-of-rekognition.

[11] Washington Post (Jay Greene), Microsoft won't sell police its facial-recognition technology, following similar moves by amazon and ibm, 2020. URL: https://www.washingtonpost.com/technology/2020/06/11/microsoft-facial-recognition/.

[12] J. Buolamwini, T. Gebru, Gender shades: Intersectional accuracy disparities in commercial gender classification, in: S. A. Friedler, C. Wilson (Eds.), Proceedings of the 1st Conference on Fairness, Accountability and Transparency, volume 81 of *Proceedings of Machine Learning Research*, PMLR, New York, NY, USA, 2018, pp. 77–91. URL: http://proceedings.mlr.press/v81/buolamwini18a.html.

[13] BBC (David Molloy), The great graphics card shortage of 2020 (and 2021), BBC, 2020. URL: https://www.bbc.com/news/technology-55755820.

[14] T. Dettmers, Which gpu(s) to get for deep learning: My experience and advice for using gpus in deep learning, 2020. URL: https://timdettmers.com/2020/09/07/which-gpu-for-deep-learning/.

[15] NVIDIA, Geforce rtx 3080, 2020. URL: https://www.nvidia.com/en-us/geforce/graphics-cards/30-series/rtx-3080/.

[16] ASUS, Recommended power supply unit table, 2020. URL: https://dlcdnets.asus.com/pub/ASUS/

Accessory/Power_Supply/Manual/RECOMMENDED_PSU_TABLE.pdf.

[17] S. Luccioni, Visualizing climate impact with gans w/ sasha luccioni, TWIML AI Podcast with Sam Charrington, 2020. URL: https://twimlai.com/visualizing-climate-impact-with-gans-w-sasha-luccioni/.

[18] C. Hill, R. Bellamy, T. Erickson, M. Burnett, Trials and tribulations of developers of intelligent systems: A field study, in: 2016 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC), 2016, pp. 162–170. doi:10.1109/VLHCC.2016.7739680.

[19] O. Vinyals, D. Silver, et al., AlphaStar: Mastering the Real-Time Strategy Game StarCraft II, https://deepmind.com/blog/article/alphastar-mastering-real-time-strategy-game-starcraft-ii, 2019.

[20] O. Vinyals, I. Babuschkin, W. M. Czarnecki, M. Mathieu, A. Dudzik, J. Chung, D. H. Choi, R. Powell, T. Ewalds, P. Georgiev, J. Oh, D. Horgan, M. Kroiss, I. Danihelka, A. Huang, L. Sifre, T. Cai, J. P. Agapiou, M. Jaderberg, A. S. Vezhnevets, R. Leblond, T. Pohlen, V. Dalibard, D. Budden, Y. Sulsky, J. Molloy, T. L. Paine, C. Gulcehre, Z. Wang, T. Pfaff, Y. Wu, R. Ring, D. Yogatama, D. Wünsch, K. McKinney, O. Smith, T. Schaul, T. Lillicrap, K. Kavukcuoglu, D. Hassabis, C. Apps, D. Silver, Grandmaster level in starcraft ii using multi-agent reinforcement learning, Nature 575 (2019) 350–354. URL: https://doi.org/10.1038/s41586-019-1724-z. doi:10.1038/s41586-019-1724-z.

[21] Next Platform (Paul Teich), Tearing apart google's tpu 3.0 ai coprocessor, 2020. URL: https://www.nextplatform.com/2018/05/10/tearing-apart-googles-tpu-3-0-ai-coprocessor/.

[22] Press Enterprise (Geri Gibbons), Talen plans data warehouse, Press Enterprise, 2021. URL: https://www.pressenterpriseonline.com/daily/021721/page/1/story/talen-plans-data-warehouse.

[23] Y. A. Strengers, Designing eco-feedback systems for everyday life, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '11, Association for Computing Machinery, New York, NY, USA, 2011, p. 2135–2144. URL: https://doi.org/10.1145/1978942.1979252. doi:10.1145/1978942.1979252.

[24] Pew Research Center (Juliana Menasce Horowitz and Ruth Igielnik and Rakesh Kochhar), Most americans say there is too much economic inequality in the u.s., but fewer than half call it a top priority, 2020. URL: https://www.pewsocialtrends.org/2020/01/09/views-of-economic-inequality/.

[25] A. W. Bartik, M. Bertrand, Z. Cullen, E. L. Glaeser, M. Luca, C. Stanton, The impact of covid-19 on small business outcomes and expectations, Proceedings of the National Academy of Sciences 117 (2020) 17656–17666. URL: https://www.pnas.org/content/117/30/17656. doi:10.1073/pnas.2006991117.

[26] E. Saez, G. Zucman, Trends in US Income and Wealth Inequality: Revising After the Revisionists, Technical Report, National Bureau of Economic Research, 2020. URL: https://www.nber.org/papers/w27921.

[27] A. Madsen, A. R. Johansen, Neural arithmetic units, in: International Conference on Learning Representations, 2020. URL:

https://openreview.net/forum?id=H1gNOeHKPS.

[28] A. Madsen, Neural Arithmetic Units and Experiences as an Independent ML Researcher with Andreas Madsen, TWIML AI Podcast with Sam Charrington, 2020. URL: https://twimlai.com/twiml-talk-382-neural-arithmetic-units-experiences-as-an-independent-ml-researcher-with-andreas-madsen/.

[29] J. Dodge, Q. V. Liao, Y. Zhang, R. K. E. Bellamy, C. Dugan, Explaining models: An empirical study of how explanations impact fairness judgment, in: Proceedings of the 24th International Conference on Intelligent User Interfaces, IUI '19, ACM, New York, NY, USA, 2019, pp. 275–285. URL: http://doi.acm.org/10.1145/3301275.3302310. doi:10.1145/3301275.3302310.

[30] U. Ehsan, M. O. Riedl, Human-centered explainable ai: Towards a reflective sociotechnical approach, 2020. URL: https://arxiv.org/abs/2002.01092.

[31] U. Ehsan, Q. V. Liao, M. Muller, M. O. Riedl, J. D. Weisz, Expanding explainability: Towards social transparency in ai systems (2021). URL: https://arxiv.org/abs/2101.04719.

[32] M. Chromik, M. Eiband, S. T. Völkel, D. Buschek, Dark patterns of explainability, transparency, and user control for intelligent systems, in: IUI Workshops, 2019.

[33] United States v. International Business Machines Corp, 1996. URL: https://www.law.cornell.edu/supct/html/95-591.ZO.html.

[34] United States v. Microsoft Corporation, 2018. URL: https://www.oyez.org/cases/2017/17-2.

[35] US and Plaintiff States v. Google, LLC, 2020. URL: https://www.justice.gov/atr/case/us-and-plaintiff-states-v-google-llc.

[36] C-SPAN, Committees of the 116th congress: House judiciary subcommittee on antitrust, commercial and administrative law, 2020. URL: https://www.c-span.org/congress/committee/?112051Źcongress=116.

[37] Washington Post (Tony Romm), U.s., states sue facebook as an illegal monopoly, setting stage for potential breakup, 2020. URL: https://www.washingtonpost.com/technology/2020/12/09/facebook-antitrust-lawsuit/.

[38] M. I. Jordan, What are the implications of algorithmic thinking? with michael i. jordan, TWIML AI Podcast with Sam Charrington, 2020. URL: https://twimlai.com/what-are-the-implications-of-algorithmic-thinking-with-michael-i-jordan/.

[39] C. Fry, H. Lieberman, Why Can't We All Just Get Along? How Science Can Enable A More Cooperative Future, Ingram Content Group, 2018. URL: https://www.whycantwe.org/.

[40] Wired (Satya Nadella and Marco Iansiti), Want a more equitable future? empower citizen developers, 2020. URL: https://www.wired.com/story/want-a-more-equitable-future-empower-citizen-developers/.

[41] D. Marculsecu, The case for hardware-ml model co-design with diana marculescu, TWIML AI Podcast with Sam Charrington, 2020. URL: https://twimlai.com/twiml-talk-391-the-case-for-hardware-ml-model-co-designwith-diana-marculescu/.

[42] B. E. Bejnordi, T. Blankevoort, M. Welling, Batch-shaped channel gated networks, CoRR abs/1907.06627 (2019). URL: http://arxiv.org/abs/1907.06627.

[43] E. Strubell, A. Ganesh, A. McCallum, Energy and policy considerations for deep learning in NLP, in: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, Florence, Italy, 2019, pp. 3645–3650. URL: https://www.aclweb.org/anthology/P19-1355. doi:10.18653/v1/P19-1355.

[44] E. Cai, D.-C. Juan, D. Stamoulis, D. Marculescu, *NeuralPower*: Predict and deploy energy-efficient convolutional neural networks, in: M.-L. Zhang, Y.-K. Noh (Eds.), Proceedings of the Ninth Asian Conference on Machine Learning, volume 77 of *Proceedings of Machine Learning Research*, PMLR, 2017, pp. 622–637. URL: http://proceedings.mlr.press/v77/cai17a.html.

[45] A. Lacoste, A. Luccioni, V. Schmidt, T. Dandres, Quantifying the carbon emissions of machine learning, CoRR abs/1910.09700 (2019). URL: http://arxiv.org/abs/1910.09700.

[46] D. Rolnick, P. L. Donti, L. H. Kaack, K. Kochanski, A. Lacoste, K. Sankaran, A. S. Ross, N. Milojevic-Dupont, N. Jaques, A. Waldman-Brown, A. Luccioni, T. Maharaj, E. D. Sherwin, S. K. Mukkavilli, K. P. Körding, C. Gomes, A. Y. Ng, D. Hassabis, J. C. Platt, F. Creutzig, J. Chayes, Y. Bengio, Tackling climate change with machine learning, CoRR abs/1906.05433 (2019). URL: http://arxiv.org/abs/1906.05433.

[47] K. Vu, Two things you need to know about reinforcement learning – computational efficiency and sample efficiency, 2020. URL: https://www.kdnuggets.com/2020/04/2-things-reinforcement-learning.html.

[48] P. Henderson, J. Hu, J. Romoff, E. Brunskill, D. Jurafsky, J. Pineau, Towards the systematic reporting of the energy and carbon footprints of machine learning, Journal of Machine Learning Research 21 (2020) 1–43. URL: http://jmlr.org/papers/v21/20-312.html.

[49] ACM Code 2018 Task Force, Acm code of ethics, 2018. URL: https://www.acm.org/code-of-ethics.

[50] K. Holstein, J. Wortman Vaughan, H. Daumé, M. Dudik, H. Wallach, Improving fairness in machine learning systems: What do industry practitioners need?, in: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI '19, Association for Computing Machinery, New York, NY, USA, 2019, p. 1–16. URL: https://doi.org/10.1145/3290605.3300830. doi:10.1145/3290605.3300830.