Senator Shoebridge

Google's language has changed from ambitious to get to the 2030 target to I think the latest phrase is "extremely ambitious". If I quote from what Google said, it said "the uncertainty around the future environmental impact of AI, which is complex and difficult to predict, is the reason why the 2030 target went from 'ambitious' to 'extremely ambitious'." Can you give us some transparency on just how much energy is being used by the expansion of AI?

Google has a long-standing commitment to sustainability, guided by our ambitious sustainability goals—including our goal of achieving net-zero emissions by 2030. We remain dedicated to these goals, and to developing AI responsibly and minimising its environmental impact.

While data centres consume electricity and contribute to emissions, cloud and hyperscale data centres collectively represent only an estimated 0.1–0.2% of global GHG emissions, based on the most recent global estimates as of 2022. While only a portion of these are accounted for by AI, the rapid growth in AI applications, especially with the impact of generative AI, is expected to drive increases going forward.

While our total emissions in 2023 increased at a slower rate compared to the previous two years, we still expect our total emissions to continue to rise before dropping toward our absolute emissions reduction target, as our business and industry continue to evolve.

We have teams across the company working harder than ever to tackle net zero. For example, in terms of AI model optimisation, Google has long been at the forefront of AI and machine learning, and we have been evolving years of deep learning research into techniques that make training faster and more efficient— enabling models that are higher quality, faster, and less compute-intensive to serve. We've identified and tested practices that our research shows can, when used together, reduce the energy required to train an AI model by up to 100 times and reduce associated emissions by up to 1,000 times, which are all used at Google today. We've sped up AI model training through techniques like quantisation, boosting large-language model training efficiency by 39% on Cloud TPU v5e. In addition, our Go Green Software guide helps developers reduce their digital footprints.

Ultimately this isn't just about Google—achieving a net-zero future for all will require systems-level change, strong government policies, and new technologies to deliver at scale this decade.

The Pawsey supercomputer in Perth is increasingly allocating, rather than time on their computer, they're allocating a certain amount of energy to researchers. Is Google doing any of that kind of mapping?

Google data centres aren't supercomputers in the same sense as Pawsey (rather a distributed network of computers), so we don't "allocate" time or energy as we might have many workloads running in complex configurations all at once. We track usage per workload and product carefully in various dimensions, including energy consumption.

In terms of energy, Google is planning to be carbon neutral by 2030 and our data centre operations form a large part of our energy footprint. This whitepaper explains the science behind our 24/7 carbon-free energy goal and how we plan to achieve it.

Is Google thinking that as a responsible global player, some kind of in-real-time acknowledgement of the energy use, feedback and response might be useful?

Google data centres are some of the most energy efficient in the world. Since 2017, we've matched 100% of the electricity consumption of our global operations with purchases of renewable energy, on an annual basis. But we're not stopping there, and aim to be carbon neutral by 2030. Google publishes a public tracker for the carbon efficiency of data centre energy consumption, which you can find here.

In addition to this, Google provides Cloud customers the ability to measure, report, and reduce their cloud carbon emissions through our Carbon Footprint tool, allowing customers to:

- View the location-based and market-based emissions derived from their Google Cloud usage, providing transparency into emissions associated with their cloud applications.
- Monitor their cloud emissions over time by project, product, and region—giving IT teams and developers metrics that can help them improve their carbon footprint.

Our detailed calculation methodology is published so that reviewers and reporting teams can verify that their emissions data meets GHG Protocol. This is all part of Google's long-standing commitment to sustainability.

Could you on notice provide Google's understanding about what the relative energy use is, for say, a simple Google search, maybe using generative AI to produce a static image, compared to producing a 10 second video?

Energy use of all of these digital services varies at the time of service based on the complexity of the computation required for a given query or query chain, and over time based on the efficiency of algorithms and chips used to enable these services. This, combined with changes in user behaviour over time, is one of the reasons we focus on top line carbon neutrality as the best way to reach our environmental goals.

The frequently cited data point that the average electricity demand of a typical Google search uses 0.3 Wh of electricity was shared more than 15 years ago. Since then our business has evolved in terms of energy efficiencies and clean energy use, and we can't rely on that specific calculation.

As disclosed in a recent earnings call, a number of technical breakthroughs are enhancing machine speed and efficiency, including the new family of Gemini models and a new generation of Tensor Processing Units (TPUs).

The recently announced Trillium, our sixth-generation TPU, is the most performant and energy-efficient TPU to date and become our most sustainable with Trillium TPUs are over 67% more energy-efficient than TPU v5e.

In addition to this, since introducing generative AI in Search, machine costs associated with generative AI responses have decreased 80% from when first introduced in Labs. This has been driven by hardware, engineering and technical breakthroughs. We delivered a number of software and hardware improvements to improve efficiency, including the development of more efficient, faster and higher performing models.

We will continue to evaluate and test ways to drive compute efficiency at scale as AI functionalities continue to roll out more widely.

Can I ask you then quickly to touch upon your current experience in the United States? Whatever we're doing here - there's this current petri dish we could look at, which is the US presidential elections. What's been the scale of deepfakes? Do you have any metrics or numbers, or some idea about what potential role they are playing?

We are committed to supporting democratic processes by surfacing high-quality information to voters, safeguarding our platforms from abuse and equipping campaigns with the best-in-class security tools and training. Our teams work around the clock to support elections by helping voters find authoritative information; working with campaigns to equip them with best-in-class security features and help them connect with voters; and protecting our platforms from abuse.

To safeguard our platforms, we have long standing policies that inform how we approach areas like manipulated media, hate and harassment, incitement to violence, and demonstrably false claims that could undermine democratic processes. We report our enforcement of these policies via quarterly Transparency Reports.[7]

We also publish regular updates concerning coordinated influence operations uncovered and terminated on our platforms through work done by the Threat Analysis Group.[8]

[7] Google, *Transparency Report*
[8] Google, *Threat Analysis Group (TAG) Official Blog*