

Lessons Learned from Trying to Argue with an AI Sceptic

The episode began with a political meme I posted: Donald Trump and Benjamin Netanyahu in orange prison jumpsuits, sitting on a bunk bed beneath a warm, nostalgic Christmas overlay reading “All I Want for Christmas.” The visual irony was immediate and sharp. Creating it required deliberate workarounds. Contemporary image-generation models have both policy safeguards and technical coherence limitations:

- Grok allows caricatures of prominent figures but consistently fails at producing reliable overlaid text.
- ChatGPT excels at generating decorative festive text like “All I Want for Christmas” but its safeguards refuse prompts depicting living political leaders in prison settings.

No single model could produce the complete image. The contradictory elements — charged political satire combined with sentimental holiday messaging — trigger refusal mechanisms or coherence failures. LLMs are simply incapable of synthesising such conceptually opposed components in one coherent output. I generated the two elements separately, then manually merged and edited them in GIMP. The final composite was undeniably human-generated: my concept, my selection of components, my assembly and adjustments. Without these tools, the satire would have remained trapped in my head or emerged as crude stick figures — stripped of all visual impact.

Someone reported the image as “AI generated.” The next day, the server introduced a new rule banning generative AI content. This rule — and the meme that triggered it — directly inspired me to write and publish the essay “High-Dimensional Minds and the Serialization Burden: Why LLMs Matter for Neurodivergent Communication.” I hoped it would encourage reflection on how these tools serve as cognitive and creative accommodations. But it turned into a rather awkward exchange with the admin.

The Sceptic’s Position and the Exchange

The admin argued that LLMs are not developed for human benefit but foster resource waste and militarization. He cited energy consumption, military ties, model collapse, hallucinations, and the risk of a “dead internet.” He revealed he had only skimmed the essay and admitted owning a powerful gaming workstation capable of running advanced local LLMs for private amusement, with access to even larger models through a friend.

Several contradictions emerged:

- My work occurs on a low-power, repairable Raspberry Pi 5 (5–15 W) using shared cloud instances. His local setup consumes far more dedicated energy and hardware.

- The hardware he uses to “tinker” with powerful LLMs locally comes from companies (Intel, AMD, NVIDIA) with billions in direct DoD contracts.

Most strikingly, the person enforcing the ban to protect authenticity was dismissing someone actively stress-testing LLMs for factual and geopolitical bias (see my public audits of Grok and ChatGPT).

The Hawking Analogy and the Admin’s Own Words

The admin self-identified as neurodivergent and acknowledged the potential of AI as assistive technology. He praised real-time captioning glasses for the visually impaired as “really cool,” but insisted that “having a machine write essays and draw pictures is different.” He added: “Neurodivergent people can do these things, many have overcome barriers in order to develop these skills.” He also described his own experience with LLMs: “The more I already know about a topic, the less I need AI. The less I know on a topic, the less equipped I am to notice hallucinations and correct them.” These statements reveal a profound asymmetry in how accommodations are judged.

Imagine applying the same logic to Stephen Hawking:

“We recognise that a voice synthesiser could help you communicate more quickly, but we’d prefer you try harder with your natural voice. Many people with motor neurone disease have overcome barriers to speak clearly — you should develop those skills too. The machine is doing something different from real speech.”

Or, from his own perspective on factual accuracy:

“The more Hawking already knows about cosmology, the less he needs the synthesiser. The less he knows, the less equipped he is to notice errors in the machine voice and correct them.”

No one would accept this. We understood that Hawking’s synthesiser was not a crutch or dilution — it was the essential bridge that allowed his extraordinary mind to share its full depth without insurmountable physical barriers.

The admin’s comfort with linear, human-scaffolded prose reflects a cognitive style that aligns more closely with neurotypical expectations. My profile is the inverse: factual and logical depth comes naturally (as in developing a multilingual publishing platform entirely myself), but producing scaffolded, accessible prose for human audiences has always been the barrier — exactly what the essay describes. To accept captioning glasses or alt-text as legitimate accommodations while rejecting LLM scaffolding for cognitive divergence is to draw an arbitrary border. Mastodon and the broader Fediverse often pride themselves on inclusivity. Yet this introduces new gates: certain accommodations are welcomed; others must be overcome through individual effort.

Historical Echoes: Resistance to Transformative Tools

The blanket rejection of public generative AI use echoes a recurring pattern throughout technological history. In early-19th-century England, skilled weavers known as Luddites smashed mechanised looms that threatened their craft and livelihoods. Gas-lamp lighters in cities opposed Edison's incandescent bulb, fearing obsolescence. Coachmen, stable hands, and horse breeders resisted the automobile as an existential threat to their way of life. Professional scribes and draftsmen viewed the photocopier with alarm, believing it would devalue meticulous handwork. Typesetters and printers fought computerised composition systems.

In every case, the resistance stemmed from genuine fear: new technology made the skills they took pride in obsolete, challenging their economic roles and social identity. The changes felt like devaluation of human labour.

Yet history evaluates these innovations by their broader impact: mechanisation reduced drudgery and enabled mass production; electric lighting extended productive hours and improved safety; automobiles granted personal mobility; photocopiers democratised information access; digital typesetting made publishing faster and more accessible. Few today would revert to gas lamps or horse-drawn transport simply to preserve traditional jobs. The tools expanded human capability and participation far more than they diminished it.

Generative AI - used as a prosthesis for cognition or creativity - follows the same trajectory: it does not eradicate human intent but extends expression to those whose ideas have been constrained by execution barriers. Rejecting it outright risks repeating the Luddite impulse — defending familiar processes at the cost of broader participation.

Conclusion: Who Decides Which Accommodations Are Acceptable?

The events recounted in this essay - one reported image, one hastily imposed ban, one protracted debate—reveal more than a local disagreement over technology. They expose a far deeper and more fundamental question: **Who gets to decide which accommodations are acceptable, and which are not?** Should it be the people who live inside the skin and brain that need the accommodation - the ones who know, from daily experience, what bridges the gap between their capabilities and full participation? Or should it be outsiders, however well-intentioned, who do not share that lived reality and therefore cannot feel the weight of the barrier?

History answers this question repeatedly, and almost always in the same direction. Wheelchairs were once criticised as encouraging dependence; deaf education systems long insisted children learn reading lips and oral speech instead of sign language. In every case, the people closest to the impairment eventually prevailed - not because they denied concerns of cost, access, or potential misuse, but because they were the primary authorities on what actually restored their agency and dignity.

With large language models and other generative tools, we are living through the same cycle again. Many who gatekeep their use do not experience the specific cognitive or expressive barriers that make linear scaffolding, narrative flow, or rapid serialization feel like an exhausting foreign-language translation task. From the outside, “just try harder” or “develop the skill” can sound reasonable. From the inside, the tool is not a shortcut around effort; it is the ramp, the hearing aid, the prosthetic that finally lets pre-existing effort reach the world.

The deepest irony emerges when the arbiters self-identify as neurodivergent, yet their particular neurology aligns more closely with neurotypical expectations in the domain being judged. “I overcame it this way, so others should too” is understandable, but it still functions as gatekeeping - replicating the very norms we critique when they come from neurotypical authorities. A consistent ethical principle is overdue:

- The person closest to the impairment is the primary authority on what enables their meaningful participation.
- External critique is legitimate on collective harms (environmental impact, misinformation risk, labour displacement), but not on the internal legitimacy of the accommodation itself.

One particularly revealing double standard appears in the widespread demand that generative AI use be explicitly disclosed. We do not require similar disclosure for most other accommodations. On the contrary, we actively celebrate technological advances that make them invisible: thick glasses replaced by contact lenses or refractive surgery; bulky hearing aids miniaturised into near-invisibility; medication for focus, mood, or pain taken privately without footnote or disclaimer. In these cases, society treats discreet, hidden use as progress - as a restoration of dignity and normality. Yet when the accommodation extends cognition or expression, the script flips: now it must be flagged, announced, justified. Invisibility becomes suspicious rather than desirable. This selective demand for transparency is not truly about preventing deception; it is about preserving comfort with a particular image of unassisted human authorship. Physical corrections are permitted to vanish; corrections to the mind must remain conspicuously marked.

If we are to be consistent, we must either demand disclosure for every accommodation (an absurd and invasive requirement) or stop singling out cognitive tools for special scrutiny. The principled position - the one that respects autonomy and dignity - is to allow each person to decide how visible or invisible their accommodation should be, without punitive rules that target one form of assistance because it unsettles existing notions of creativity and intellect. This essay is not merely a defence of one particular tool. It is a defence of the broader right of disabled and neurodivergent people to define their own access needs, without having to justify them to those who have never walked in their shoes. That right should not be controversial. Yet, as the preceding account shows, it still is.