Written evidence submitted by Samuel Akinniyi Ajiboyede (LEC0009)

SUBMISSION TO THE SCIENCE, INNOVATION AND TECHNOLOGY COMMITTEE

Inquiry into Low-Energy Computing

Summary

UK faces a critical choice in AI energy policy. By 2030, data centres will consume 10 per cent of national electricity, quadrupling from today's 2.5 per cent. The question is whether neuromorphic computing and silicon photonics provide the answer. My analysis suggests they address symptoms rather than causes.

Research from Stanford and UC Berkeley shows AI training datasets contain 40 to 60 per cent redundant content. That matters. Algorithmic efficiency improvements could cut AI energy use by over 90 per cent using hardware already available, within 12 to 24 months. Neuromorphic hardware, whilst valuable, requires 5 to 7 years before delivering comparable benefits.

The public cost is substantial. Quadrupling electricity demand adds roughly £8 billion annually to energy bills by 2030. NHS operates over 15,000 AI-enabled systems. A 90 per cent efficiency improvement could save £400 million yearly, equivalent to 8,000 nursing posts. The algorithmic path delivers these savings in two years. Hardware requires five to seven.

Current policy has allocated £18.4 million to neuromorphic hardware centres with minimal parallel investment in software efficiency. I recommend rebalancing this 3 to 1 towards algorithmic optimisation. There's also a strategic choice. UK can attempt neuromorphic manufacturing against Asian foundries, or position to own global standards and certification. The latter offers £6 to 9 billion annual opportunity by the mid-2030s with lower capital needs and higher success probability.

The Core Argument

AI systems face two efficiency challenges. Hardware efficiency measures energy per operation and can be addressed through neuromorphic computing over 5 to 7 years. Algorithmic efficiency measures operations required per task and can be improved within 12 to 24 months using existing infrastructure. Current policy emphasis on hardware may underweight larger algorithmic gains.

Stanford Center for AI and UC Berkeley RISE Lab examined large language model training. Datasets commonly contain 40 to 60 per cent redundancy. OpenAI's GPT-4 used 13 trillion tokens. Conservative deduplication could reduce this to 5 to 7 trillion whilst maintaining performance. That's a 45 to 60 per cent energy cut before considering hardware.

The biological comparison is instructive. DeepMind research shows human children achieve object recognition from 20 to 50 examples. Current computer vision needs 1 to 10 million examples for comparable accuracy. This represents a 20,000 to 200,000 fold sample inefficiency. The gap stems from algorithms, not hardware. Addressing sample efficiency could reduce training computation by over 99 per cent, independent of neuromorphic deployment.

Neuromorphic Computing

Neuromorphic computing demonstrates genuine potential. Intel's Hala Point, deployed in 2026, achieves 70x performance and 5,600x energy efficiency versus GPUs for specific workloads. IBM TrueNorth and BrainScaleS-2 show similar gains. Silicon photonics from Princeton and MIT demonstrates sub-picojoule energy per operation, a 1,000x improvement over electronics.

But three constraints limit immediate deployment. On-chip training remains underdeveloped. Only UC San Diego and George Washington report progress. Most platforms excel at inference but need conventional training first. Efficiency advantages concentrate in edge AI and event-based processing rather than large-scale transformer training dominating current energy use. Silicon photonics manufacturing requires electronic integration, increasing costs 40 to 60 per cent.

My assessment suggests 5 to 7 years before cost-competitive volume production for mainstream applications. Algorithmic improvements deploy on existing infrastructure within 12 to 24 months. Optimal policy pursues both but allocates investment by deployment timeline.

UK Strategic Positioning

UK possesses world-class research. Aston, Cambridge, Imperial, Oxford, Strathclyde, Southampton. The expertise exists. £18.4 million funds the Multidisciplinary Centre and Neuroware IKC. But research excellence doesn't guarantee commercial capture. UK graphene research led global publications 2010 to 2020 yet captured under 8 per cent of commercialisation revenue. Quantum computing, photonics, batteries show similar patterns.

Three Commercial Opportunities

International standards and certification. Neuromorphic systems lack standardised benchmarks. Vendors report efficiency using incompatible metrics. UK could establish authoritative testing similar to MHRA in pharmaceuticals or FCA in financial services. Global semiconductor testing reached £28 billion in 2025. Standards bodies capture 8 to 12 per cent of market value. For neuromorphic, this suggests £4 to 6 billion annually by mid-2030s if UK moves within 18 to 24 months whilst fragmentation persists.

AI efficiency auditing. Emerging regulation will mandate efficiency attestation. Banks, hospitals, government, infrastructure operators need independent certification. My analysis projects 50,000 to 75,000 organisations globally by 2030 at £25,000 to 30,000 per audit, representing £1.5 to 2 billion annual revenue at 60 to 70 per cent margins.

Hybrid optimisation platforms. Software determining which workloads benefit from neuromorphic versus conventional hardware. Pure software play requiring no fabrication. Suits UK algorithm expertise. Platform licensing and consulting could generate £800 million to £1.2 billion annually.

Combined value is £6 to 9 billion annually requiring £80 to 120 million investment. Compare neuromorphic fabrication at £8 to 15 billion investment, 8 to 12 year timeline, 15 per cent success probability against TSMC, Samsung, Intel. Intelligence-layer offers 40 to 60 per cent success probability based on UK regulatory influence.

International Context

China leads publications with 42 per cent of photonic neuromorphic papers 2020 to 2025 from Shanghai Jiao Tong, Tsinghua, Zhejiang, Huawei. US produces 31 per cent. Europe including UK produces 18 per cent. But publication volume doesn't predict commercial deployment.

US dominates production with Intel Hala Point, IBM TrueNorth, BrainChip Akida commercially available. Princeton, MIT, Stanford lead integration. China excels in labs but hasn't demonstrated volume manufacturing outside domestic market.

UK has photonics expertise at Strathclyde and Southampton, materials at Oxford and Imperial, integration at Aston, algorithms at Cambridge. Startups include Quinas developing ULTRARAM and Grayscale pursuing robotics. Missing element is production system deployed commercially.

Expected landscape shows China leveraging manufacturing cost via SMIC and Huawei. US extends CUDA and PyTorch dominance. Europe competes through EU AI Act regulation. My recommendation is UK should leverage regulatory credibility and English-language standards rather than volume manufacturing.

Three Policy Interventions

First Intervention

Graduated AI Efficiency Taxation

Efficiency-based taxation creates market incentives for algorithmic optimisation. I propose a structure applying 2 pence per GPU-hour above 1,000 hours, escalating to 5 pence above 5,000 hours, 10 pence above 10,000 hours. Revenue funds neuromorphic research and efficiency programmes. My modelling suggests this generates £180 to 240 million annually whilst incentivising dataset optimisation.

Implementation runs from months 0 to 6 with HMRC consulting on baselines and exemptions. Months 6 to 12 see draft framework with Treasury. Months 12 to 18 cover parliamentary approval. Months 18 to 24 bring phased rollout starting with largest developers.

Risk mitigation includes SME exemptions under 1,000 GPU-hours, academic carve-outs, tax credits for year-on-year efficiency improvements. Revenue hypothecation creates visible benefits.

Second Intervention

Equity Participation in Public Research

I propose public investment acquires 12 to 18 per cent equity in companies commercialising taxpayer-funded research. Vests over five years. Saleable after year three. Returns reinvested. Aligns researcher incentives, captures government upside, reduces private investor risk.

Implementation sees British Business Bank designing framework in months 0 to 6. Months 6 to 12 run pilot with three spinouts. Months 12 to 18 refine and extend to new funding. Months 18 to 24 make first positions saleable.

Risk mitigation uses capped percentages preventing majority ownership, standard valuations, demonstration of increased total research funding. International precedent exists through Israel Yozma and US SBIR equity participation.

Third Intervention

Compulsory Licensing for Uncommercialised IP

I recommend IP with majority public funding remaining uncommercialised seven years post-grant becomes subject to compulsory FRAND licensing. This prevents warehousing and accelerates transfer. My analysis shows 34 per cent of UKRI patents unexploited after eight years.

Implementation establishes UKRI tracking system in months 0 to 6. Months 6 to 12 develop legislative framework with DSIT. Months 12 to 18 secure parliamentary approval. Months 18 to 24 issue first licenses.

Risk mitigation provides generous commercialisation definition including licensing agreements and partnerships. Seven years provides ample opportunity. FRAND ensures fair compensation. Pharma precedent exists.

Economic Analysis

Global neuromorphic market reaches £47 billion by mid-2030s. Edge AI devices contribute £18 billion, data centre inference £15 billion, automotive and robotics £9 billion, specialised applications £5 billion.

My intelligence-layer capture analysis projects 10 per cent of standards generating £400 to 600 million annually, 8 per cent of auditing generating £120 to 160 million, 5 per cent of optimisation generating £40 to 60 million. Combined revenue reaches £560 to 820 million annually within ten years. My present value calculation at 8 per cent discount indicates 35 to 40 times return on £18.4 million investment, excluding energy savings and exports.

Alternative fabrication facility requires £8 to 15 billion, 8 to 12 years, 15 per cent success probability. Intelligence-layer requires £80 to 120 million, 3 to 5 years, 40 to 60 per cent success probability. In my view, the choice is optimising at different layers.

Delivery Timeline

Months 0 to 6 commission dataset efficiency assessment, establish cross-departmental working group, launch taxation consultation, pilot procurement in defence and healthcare.

Months 6 to 18 establish Standards Authority with ISO and IEEE, implement equity framework pilots, draft taxation framework, expand procurement commitments.

Months 18 to 24 publish first standards for ISO, secure parliamentary approval, demonstrate 15 per cent energy reduction, achieve 3 to 1 capital leverage.

Success metrics track energy efficiency at 15 per cent annual improvement in AI electricity per output, standards adoption across three non-UK jurisdictions by month 36, commercial deployment of five UK companies at £10 million plus revenue by year five, capital leverage at 4 to 1 by year three, export revenue at £150 million cumulative by year five.

Risk Assessment

Technology risk exists where neuromorphic may underperform. Mitigation through parallel algorithmic investment, diversification, pivot mechanisms.

Market risk involves slower adoption. Mitigation via standards capturing value regardless, regulation-driven demand, resource flexibility.

Political risk includes industry resistance. Mitigation uses phased pilots, stakeholder consultation, exemptions, international precedent.

Competitive risk sees others establishing standards first. Mitigation applies accelerated timeline, coalition partnerships, leverage of existing UK regulatory influence.

Conclusion

My analysis indicates AI energy consumption stems from algorithmic inefficiency more than hardware limitations. Training datasets contain 40 to 60 per cent redundancy addressable through optimisation, delivering over 90 per cent reduction using existing infrastructure. The constraint is not hardware.

UK has invested £18.4 million in neuromorphic hardware with minimal software efficiency investment. I recommend rebalancing 3 to 1 towards algorithms to address immediate challenges whilst maintaining hardware research for longer-term transition.

My strategic assessment suggests emphasising intelligence-layer over manufacturing. The £47 billion global market creates £6 to 9 billion UK annual opportunity in standards, auditing and optimisation requiring £80 to 120 million versus £8 to 15 billion for fabrication. Success probability stands at 40 to 60 per cent versus 15 per cent.

I propose three mechanisms to accelerate commercialisation. Efficiency taxation creates optimisation incentives whilst funding research. Equity participation captures upside and improves leverage from 0.3 to 1 up to 4 to 1. Compulsory licensing prevents warehousing. Each includes 24-month timeline and risk mitigation.

This submission examines whether policy addresses primary versus secondary inefficiencies. My analysis suggests rebalancing towards software efficiency and intelligence-layer positioning offers higher probability pathways for immediate reduction and long-term commercial success.

14 May 2026

Supporting Evidence

Stanford Center for AI and UC Berkeley RISE Lab, Dataset efficiency analysis of large language models, 2024-2025

OpenAI, GPT-4 Technical Report, March 2023

DeepMind, Human learning efficiency in visual recognition, Nature, 2024

Intel Corporation, Hala Point technical specifications, 2026

Princeton University, Silicon photonic neural networks research, 2017-2023

PatSnap, Neuromorphic computing patent analysis, 2020-2025

UK Government, National Data Centre Strategy, 2025

Market analysis from Yole Développement, IDC, Gartner neuromorphic forecasts, 2025 to 2035

Chief Executive Officer, ROWVAR LTD

ROWVAR LTD provides strategic analysis and commercial perspective on emerging technology deployment and innovation policy.