Written Evidence Submitted by the UK Multidisciplinary Centre for Neuromorphic Computing (LEC0016)

Aston University, Birmingham

Science, Innovation and Technology Committee

Inquiry: Low-Energy Computing

About NeuroSYNC

NeuroSYNC – the UK Multidisciplinary Centre for Neuromorphic Systems and Computing – was established in 2025 with UKRI/EPSRC funding and is led from Aston University in consortium with the Universities of Oxford, Cambridge, Southampton, Strathclyde, Queen Mary, and Loughborough. The Centre is supported by 22 industrial and non-academic partners spanning defence, healthcare, telecommunications, and advanced manufacturing.

NeuroSYNC's mission is to become a sustainable focal point for neuromorphic computing research in the UK and to position the UK as a global leader in neuromorphic computing through coordinated activity across algorithms, hardware, materials, photonics, and applications. The Centre leads national road mapping activities, engages a growing community of over 100 stakeholders through structured surveys and workshops, and trains the next generation of researchers in this field.

NeuroSYNC operates within a complementary national ecosystem alongside NeuMat (a network focused on neuromorphic materials and devices) and NeuroWare (an Innovate UK Knowledge Centre pushing research toward commercialisation and spin-out). Together these three bodies are building the connective infrastructure between fundamental research, hardware development, and commercial deployment.

Executive Summary

The energy demands of AI are becoming a real and urgent strategic challenge. UK data centres currently consume approximately 2.5% of national electricity, with projections suggesting a four-fold increase by 2030. Current GPU-based computing architectures are unlikely to address this trajectory through incremental efficiency improvements alone. Neuromorphic computing and neuromorphic photonics offer a fundamentally different approach to computation — one in which the UK is scientifically and technologically well positioned to lead.

NeuroSYNC's recent stakeholder survey (100+ participants) and roadmapping workshop found that:

- 61% of respondents expect neuromorphic computing to deliver real-world impact within **three to seven years**.
- Edge AI and IoT sensing – the application domain most constrained by current energy and latency limitations – was identified as the top priority sector by 79% of respondents.

- The UK has genuine scientific strengths in this field, but requires coordinated policy action to convert research into sovereign industrial capability.

This submission addresses the Committee's specific areas of interest and makes three concrete recommendations: a national neuromorphic testbed, application-led funding instruments with longer horizons, and investment in a shared software and benchmarking stack.

Q1. How realistic is low-energy computing as a solution to AI's energy demands, and what are its current limitations?

The energy case is quantitatively strong for specific workload classes

The energy advantage of neuromorphic approaches over conventional GPU-based systems is measurable and real, but workload-specific. A high-end training GPU (e.g. NVIDIA H100) draws approximately 700 W under load. Neuromorphic chips, which emulate the brain's event-driven, sparse-activity regime, operate in a radically different envelope: IBM's TrueNorth runs at 65-70 mW at real-time throughput; Intel's Loihi platforms have reported energy-per-inference gains of 100-1,000× on sparse workloads such as keyword spotting, constraint optimisation, and graph search. Published neuromorphic targets sit in the 10-100 TOPS/W range versus 1-10 TOPS/W for current AI accelerators.

This advantage is most acute – and most actionable in the near term – in:

- Edge AI and IoT sensing: continuous low-power inference at the device, without cloud connectivity
- Real-time sensor processing: autonomous transport, robotics, defence systems where latency and power constraints are mission-critical
- Combinatorial optimisation: logistics, circuit design, and scheduling tasks where sparse, event-driven processing is inherently suited

Photonics adds a second, complementary dimension

Neuromorphic photonics – combining silicon photonics with neuromorphic computing principles – offers additional capabilities beyond electronic neuromorphic approaches alone. Recent research has demonstrated all-optical recurrent neural networks operating at 100 GHz clock rates[1], exceeding the approximately 5 GHz limit of electronic systems. An optical neural network can, in principle, process information using less than one photon per multiplication operation[2]. Light also provides inherent parallelism and bandwidth density that electronics cannot match for certain matrix computation tasks central to AI inference.

The NeuroSYNC community's assessment of photonics identifies its clearest near-term advantages as:

- High bandwidth and speed for linear operations (matrix-vector multiplication in the optical domain)

---

[1] https://arxiv.org/abs/2501.05756
[2] https://doi.org/10.1038/s41467-021-27774-8

- Sensing and in-sensor computing, exploiting optical sensitivity without requiring active laser sources
- Low-latency inference in telecom and data-centre interconnect contexts, where photonic infrastructure is already deployed

The community consensus is that photonics is most effective as part of a hybrid system – optics providing bandwidth and fast linear computation, electronics providing non-linearity and control – rather than as a standalone replacement for electronics.

Honest limitations the Committee should understand

NeuroSYNC's stakeholder community is clear-eyed about the current state:

- Energy advantages are workload-specific. They are not yet matched for the dense matrix multiplications that dominate large language model training. Neuromorphic computing is a precision instrument for a growing class of applications, not a universal replacement for GPU infrastructure.
- Hardware variability and reliability are the top integration bottleneck, cited by participants across our survey and workshop. Scaling from research prototypes to manufacturable, deployable systems remains a significant engineering challenge.
- The field lacks a widely accepted 'killer application' analogous to image classification for deep learning. One of our workshop participants captured the risk plainly: the single biggest challenge is that neuromorphic computing remains permanently 'ten years away' – impressive in research but never arriving in deployment.
- The software stack is immature. There is no CUDA equivalent for neuromorphic hardware; adoption is currently confined to specialists and cannot scale without accessible, standardised tools and benchmarks.
- Wall-plug efficiency at the system level must not be conflated with device-level efficiency. A photonic chip with excellent energy-per-operation performance may still require substantial supporting infrastructure. Benchmarks must account for full system energy, including training as well as inference.

Q2. What is the scientific research base underpinning these technologies, and how strong is UK capability?

UK scientific foundations are genuinely strong

The UK neuromorphic and photonics research community has internationally recognised strengths across the full technology stack:

- **Algorithms:** spiking neural networks, local learning rules, reservoir computing, and surrogate gradient methods, with active contributions to the development of biologically plausible learning beyond backpropagation
- **Hardware:** electronic neuromorphic platforms (including contributions to SpiNNaker, hosted at Manchester), photonic integrated circuits, and phase-change material (PCM) synaptic devices integrated with silicon photonics
- **Photonics:** world-leading capability in silicon photonics, compound semiconductors, and optical reservoir computing, within the NeuroSYNC consortium

- **Materials:** PCM research (including antimonide-selenide materials demonstrating 100+ million cycles of optical switching), with growing interest in ferroelectric materials and 2D materials for neuromorphic devices

NeuroSYNC's survey found that respondents placed UK capability across the spectrum from 'leader' to 'follower' to 'emerging', reflecting genuine research strength alongside real questions about translation and scale. The honest assessment is: strong science, insufficient translation infrastructure.

Community priorities identified through structured engagement

NeuroSYNC's roadmapping workshop – drawing together neuroscientists, algorithm researchers, hardware engineers, materials scientists, photonics specialists, and industry representatives – identified the following near-term UK research priorities:

- **Spiking neural networks (digital implementations):** supported by 61% of survey respondents; the technology most credibly deployable within five years using existing fabrication ecosystems
- **Local and unsupervised learning:** identified as strategically important for enabling on-device, adaptive systems without requiring costly retraining cycles
- **Reservoir computing:** valued for simplicity and implementability, particularly in photonic substrates; potentially the approach most amenable to a near-term 'killer application' demonstration
- **Dendritic computation and multi-scale neural modelling:** identified as a significant knowledge gap where closer neuroscience-engineering collaboration could unlock new architectural paradigms
- **Explainability:** a cross-cutting gap; neuromorphic systems remain difficult to interpret, limiting adoption in regulated sectors

A neuroscience-to-engineering bridge is needed

A recurring theme across NeuroSYNC's community engagement was the fragmentation of language between neuroscience, algorithms, hardware engineering, and policy. Researchers consistently identified the need for clearer shared definitions – of learning, plasticity, neural circuits, and the appropriate level of biological abstraction – as a precondition for productive hardware and algorithm co-design. The field's approach should be neuromorphic (inspired by biology) rather than neuromimetic (replicating it), with neuroscience providing computational principles rather than blueprints.

Specific knowledge gaps at the neuroscience interface include: the computational role of spike timing and energy efficiency; mechanisms of long- and short-term memory; dendritic integration and its implications for on-chip learning; and the role of noise in biological computation.

Q3. Does the UK have sovereign capabilities in low-energy computing, and what is needed to strengthen them?

The UK has the foundations – but translation infrastructure is absent

The UK's sovereign position in neuromorphic and photonic computing rests on three strengths: internationally competitive research groups across multiple universities; an emerging national ecosystem coordinating activity across the research-to-commercialisation pipeline; and a strong industrial partner base spanning defence, healthcare, and advanced manufacturing.

The NeuroSYNC ecosystem – NeuroSYNC (fundamental and applied research), NeuMat (materials and devices), and NeuroWare (commercialisation and spin-out) – provides the coordinating infrastructure that frontier technology demands. NeuroSYNC's 22 industrial partners include organisations from Airbus, Thales, Leonardo, QinetiQ, and BAE Systems to NHS trusts and regional SMEs. Regional engagement in the West Midlands alone has reached over 500 attendees at AI and neuromorphic events for SMEs, with 70 AI adoption roadmaps developed in partnership with regional businesses.

Critical gaps in sovereign capability

Despite these strengths, the UK lacks several components necessary for genuine sovereign capability:

- **Shared hardware access and testbed infrastructure:** Academic and SME teams currently navigate a fragmented landscape of proprietary platforms without the shared infrastructure that has underpinned progress in quantum computing. Fabrication access and supply chain dependencies represent the most critical strategic vulnerability.
- **Software stack:** The absence of a standardised, accessible software and benchmarking ecosystem means that neuromorphic hardware adoption cannot scale. Without open toolchains, simulation environments, and agreed benchmarks, the field remains confined to specialists. This is the single most consistently identified gap across NeuroSYNC's community engagement.
- **Application demonstrations at scale:** The community consistently called for working prototypes and simple, functional systems with real-world relevance – the neuromorphic equivalent of ImageNet for deep learning. Without such demonstrations, private investment and industrial adoption cannot accelerate.
- **Training benchmarks:** Current benchmarks focus on inference. Benchmarks that account for the full system energy cost of training – which may negate device-level efficiency advantages if not addressed – are a critical missing component.

The strategic argument

The UK's current AI infrastructure commitments concentrate sovereignty risk in foreign hardware and cloud providers. Neuromorphic computing is not a substitute for GPU access; it is an option on a different computing substrate, built on expertise the UK already has. The comparison to quantum computing is instructive: quantum has attracted substantial public funding, rightly, given long-term potential. But most commercially relevant quantum applications remain years from useful scale. Neuromorphic computing is closer to deployment for well-defined problem classes, and the UK is positioned to lead rather than import.

The global landscape is moving. Intel's Loihi programme, IBM's NorthPole, European initiatives through EBRAINS and BrainScaleS, and growing investment across the US and East Asia signal that neuromorphic computing is becoming an industrial competition. The question for the UK is not whether neuromorphic computing will matter, but whether the UK will be a maker or a buyer of the technology when it does.

Q4. What is the government doing to support R&D in low-energy computing, and what more should it do?

Current support: a strong foundation, but insufficiently coordinated

UKRI/EPSRC investment in NeuroSYNC, and related support for the neuromorphic materials and commercialisation ecosystem, represents a meaningful commitment. NeuroSYNC's activities – roadmapping, policy engagement, international research partnerships, early-career researcher training, and the planned International Neuromorphic Computing Conference (INCC 2027) in London – provide the coordination infrastructure that the field requires.

However, the current funding landscape has structural weaknesses that the NeuroSYNC community has consistently identified:

- Grant cycles are mismatched to the development timelines of frontier hardware. Neuromorphic hardware-software co-design requires sustained multi-year investment that does not fit neatly into standard EPSRC project structures.
- There is no equivalent to the National Quantum Computing Centre for neuromorphic computing: no shared infrastructure, no common testbed, no national facility through which the community – including SMEs – can access state-of-the-art hardware.
- The private capital needed to move from research to deployment is waiting for clearer signals. Industry partners engage actively with NeuroSYNC, but investment decisions require demonstrated hardware capability at scale and a credible route to CMOS-compatible or CMOS-integrated manufacturing.

Three recommendations for government action

Recommendation 1: Establish a National Neuromorphic Testbed

A shared, nationally accessible neuromorphic hardware facility – even modest in initial scope – would address the most critical infrastructure bottleneck. Such a facility would: accelerate translation from research to deployable capability; provide the common platform needed to develop meaningful shared benchmarks (including for training as well as inference); enable SME access to frontier hardware without requiring prohibitive capital investment; and establish a focal point for international research collaboration. The quantum testbed model provides a relevant precedent. Funding at a comparable level – matched to neuromorphic's earlier maturity stage – would be appropriate.

Recommendation 2: Application-led funding instruments with longer horizons

The NeuroSYNC community is clear that different problems require different algorithms and architectures, and that neuromorphic development does not fit neatly into standard grant cycles. Instruments that explicitly connect neuroscientists, hardware engineers, and end-user sectors – with five-to-ten year horizons and milestone-based progression – would unlock private capital currently waiting for clearer signals. These instruments should require application specificity: proposals must identify the workload, the hardware architecture, and the energy or performance metric that defines success.

Recommendation 3: Investment in the software and benchmarking stack

The single most consistent theme across all NeuroSYNC community engagement – survey, workshop, and industrial partner input – is the absence of accessible, standardised software and benchmarking tools. Without a CUDA equivalent for neuromorphic hardware, adoption will remain confined to specialists. The government should fund: open-source toolchain development aligned with the most credible near-term platforms (digital spiking neural network implementations on CMOS and FPGA); a set of agreed benchmarks covering both inference and training, covering the relevant application domains identified by the community (edge AI, robotics, defence, healthcare); and simulation environments that enable algorithm researchers without hardware access to develop and test approaches at scale. This is an area where coordinated UK investment could establish a durable competitive advantage.

Q5. How long would it take to implement low-energy computing, and when would its benefits be felt?

NeuroSYNC's stakeholder survey found that 61% of respondents expect neuromorphic computing to deliver real-world impact within three to seven years. This is neither imminent nor distant: it is a realistic horizon for well-scoped applications with the right investment.

A more granular view, based on the NeuroSYNC community's assessment:

Near-term (1-3 years): edge deployment and niche applications

- Digital spiking neural network implementations on existing CMOS and FPGA platforms are credible within this timeframe. Commercial neuromorphic processors (Intel Loihi 2, BrainChip Akida, SpiNNaker 2) exist today, and pilot deployments in edge AI, keyword detection, and sensor fusion are achievable.
- The UK's contribution in this phase is primarily algorithmic and systems-level: developing local learning rules, unsupervised adaptation, and reservoir computing approaches that exploit existing hardware.
- Photonics in this phase contributes primarily through optical interconnects and in-sensor computing, where the technology is closer to deployment and the infrastructure is partly in place.

Medium-term (3-7 years): broader deployment and UK industrial impact

- Neuromorphic hardware purpose-built for UK application priorities – defence and security, healthcare monitoring, autonomous systems – could reach deployable readiness. This requires: fabrication access (CMOS-compatible integration or dedicated neuromorphic foundry capacity); the software and benchmarking infrastructure described above; and application-level demonstrations that de-risk industrial investment.
- Photonic neuromorphic accelerators for specific matrix-computation workloads – particularly relevant to data-centre inference at scale – are achievable in this window if the integration challenges (optical-electronic interfacing, scalable manufacturing) are addressed.

- PCM-based photonic synaptic devices, which the NeuroSYNC materials community identifies as the most promising near-term neuromorphic photonics substrate (with demonstrated 100+ million switching cycles), could reach pilot-scale fabrication.

Longer-term (7+ years): systemic impact on AI energy infrastructure

- Neuromorphic computing becomes a meaningful component of the UK's frontier compute strategy – not replacing GPU infrastructure, but providing the energy-efficient, sovereign substrate for a growing proportion of AI inference workloads, particularly at the edge and in real-time sensing applications.
- The UK's 2030 clean energy targets are the relevant policy constraint. Neuromorphic computing cannot directly resolve the 2030 data-centre energy challenge: the timeline does not permit it. But investment made now determines whether the UK is a sovereign participant or a technology importer in the decade that follows.

Conclusion

NeuroSYNC submits this evidence with a clear message: the UK's position in neuromorphic computing and neuromorphic photonics is strong, the scientific foundations are in place, and the commercial case is sharpening. What is needed now is the policy ambition to match – recognising neuromorphic computing as a genuine pillar of the UK's frontier compute and clean energy strategy, and backing it with the three instruments outlined above.

The window for the UK to lead rather than follow in this technology is open. It will not remain so indefinitely.

14 May 2026

NeuroSYNC – UK Multidisciplinary Centre for Neuromorphic Systems and Computing
Aston University, Birmingham | www.neurosync.ac.uk
Contact: Centre Manager, NeuroSYNC | Aston Institute of Photonic Technologies (AIPT)