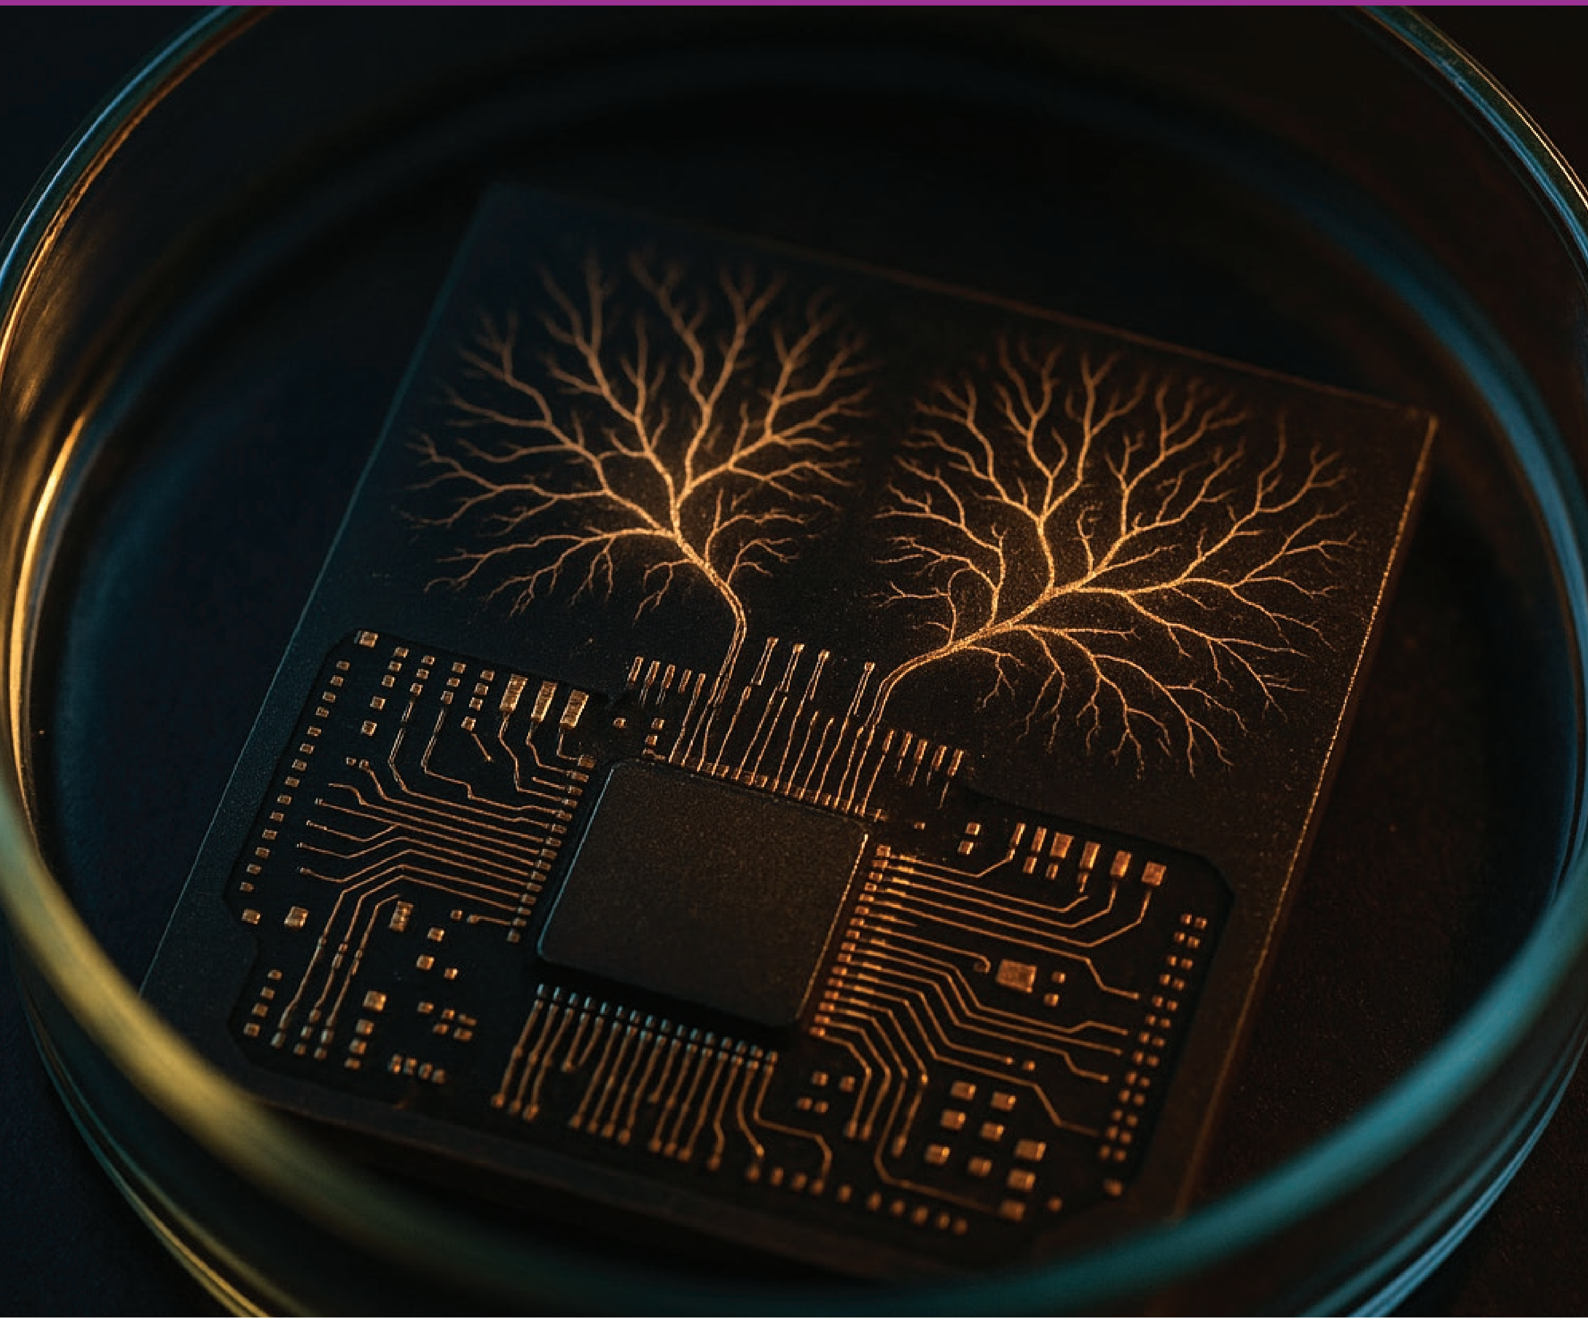

# Neuromorphic Computing

**Bill Jeffrey**
Chair of HPQC Fund Technical Advisory Board

August 2025

# Neuromorphic Computing

**Sector Overview: ICM HPQC Fund**

At the **ICM** HPQC Fund we seek out early-stage investments in technologies that can radically improve the way compute can be generated and delivered.

Beyond High Performance Classical Data Centres, AI, LLMs, Nvidia, and even quantum computing is an area that we are increasingly interested in – remember the HPQC Fund loves breakthrough technologies but only invests when those technologies are ready to scale commercially. We may be close to that tipping point with Neuromorphic Computing.

Written by the Chair of our Technical Advisory Board, Dr Bill Jeffrey, the following paper provides an introduction for our clients, an overview of the sector, and a framework for our ongoing communication on this topic with you.

Dr William "Bill" Jeffrey has worked with emerging technologies for 30+ years, driving the development of groundbreaking products and bringing innovations from the lab to the marketplace, including as Director of NIST and as CEO of HRL and SRI International. Bill is a director of unlisted and listed technology companies internationally.

> **“** Biomimicry is the conscious emulation of life’s genius. **”**
>
> **– Janine Benyus**

# Introduction

When it comes to autonomous operations and reasoning, the human brain is unsurpassed by even the most advanced AI algorithm. Despite computing speeds being much slower than silicon processors and containing a vastly smaller portion of human-generated data, the human brain can adapt quickly to radically different environments and create new mental models when confronted by a unique set of inputs.

With only 20W of power, a volume of less than 1300cm3, and a mass of ~1.4 kg, the human brain stands as existence of proof that an architecture unlike GPUs/TPUs can achieve complex reasoning in near real time with limited infrastructure. Given the efficiency of the brain, researchers and companies are developing hardware that tries to mimic the physical processes within the brain, hoping to achieve human-like intelligence. These approaches are broadly referred to as “neuromorphic computing” and this is still an emerging field.

For the most part, neuromorphic computing architectures replicate how we believe brain cells are interconnected and interact, but rather than using biological cells, they are built primarily with semiconductors. These simulated brains compute in a fundamentally different way than classical computers (referred to as von Neumann machines) and offer advantages in applications that require reasoning, adaptability, low power, and learning. Some of the major attributes of the brain are shown in Figure 1[2] (top) and the analogous neuromorphic instantiation (bottom). Many of these attributes are

**Figure 1: Top Level mapping from brains (top) to neuromorphic systems (bottom).**

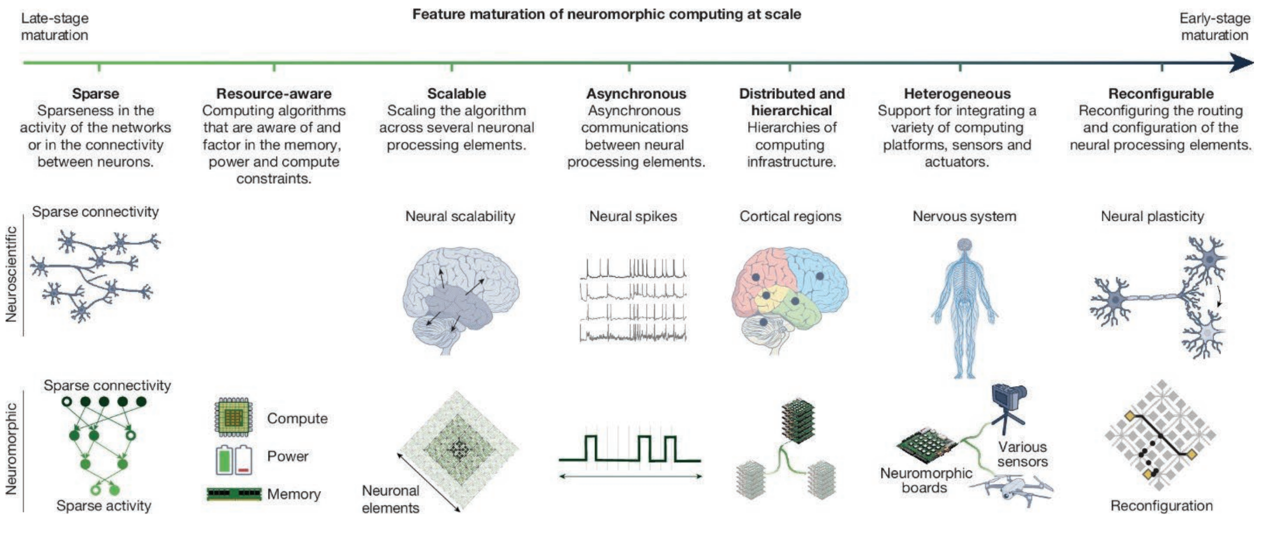


[1] See for example: https://en.wikipedia.org/wiki/Brain_size
[2] Kudithipudi, D., Schuman, C., Vineyard, C.M. et al., 2025, *Neuromorphic Computing at Scale*, **Nature**, **637**, 801.

responsible for the high efficiency of the brain, form the rationale for why neuromorphic computing is being pursued, and will be discussed in more detail below.

There are, however, impediments to adoption including the lack of standards, limited tools and experience with programming, and an immature manufacturing ecosystem. These impediments will also be discussed below. This paper will first describe how the human brain processes data, followed by neuromorphic architectures, and then delve into applications.

# How do Human Neurons Communicate and Store Data?

The fundamental building block of the human brain is a neuron (also referred to as a nerve cell – see Figure 2[3]). The neuron is electrochemically connected to other neurons via dendrites (think of them as small chemical pipes). The dendrites are branch-like extensions of a neuron that receive signals (neurotransmitters). The signals are aggregated in a neuron, and if strong enough, an electrical signal (called an action potential or spike) is generated and transmitted along an axon (analogous to a wire). At the end of each axon are multiple axon terminators, which convert the electrical signal back into a chemical signal which are received by the dendrites of neighboring neurons over a small gap (referred to as a synapse). A group of neurons working together to create a signal is called a neural circuit. Note that if the aggregated signal in the neuron is below the action potential (or spike), then no signal is transmitted.

**Figure 2: Depiction of the communication of human neurons via synapses.**

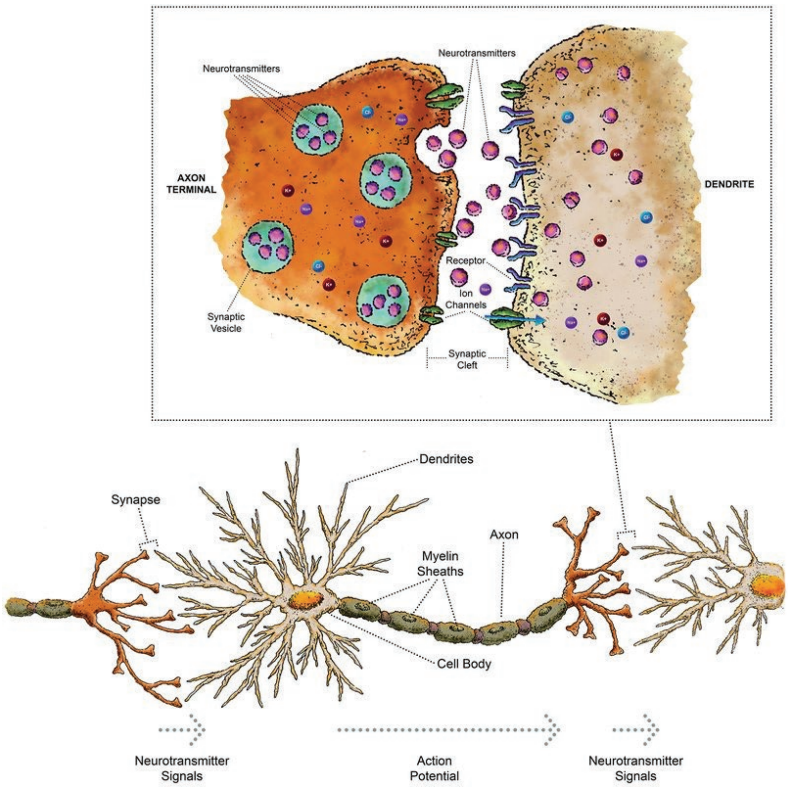


[2] Source: https://scienceexchange.caltech.edu/topics/neuroscience/neurons.

## The **human brain contains about 100 billion neurons** ($10^{11}$) with over 100 trillion connections ($10^{14}$ synapses), implying each neuron is connected to about 1,000 other neurons, creating a very complex network.

Memory is not stored within a neuron. Instead, memories are stored through changes in the synaptic connection between neurons. These connections can be either strengthened or weakened, allowing the brain to encode, store, and retrieve information. Given that an individual neuron is connected to ~1000 other neurons (via synaptic connections), a memory will be broadly distributed over many neurons – which are sharing multiple other memories. This distributed architecture provides for robustness against the inevitable damage or death of individual neurons.

Although memories are not stored in a specific neuron, the brain does compartmentalise sensory inputs and their processing into different regions of the brain[4]. For example, the occipital lobe processes visual data, the temporal lobe processes smell, taste, and sound, and the frontal lobe controls thinking, planning, organising, problem-solving, short-term memory, and movement. The ability of the brain to segregate brain regions by input stimulus is not yet captured in neuromorphic processors, and may be a fundamental gap in the hardware emulations discussed below.

Sleep appears to play a vital role in memory formation by allowing the brain to replay and stabilise memories. To retrieve a memory, the brain reactivates the neural pathways associated with that memory. Memory retrieval can be imperfect, leading to incomplete or inaccurate memories.

The strengthening and weakening of synaptic connections are fundamental properties of the brain and are referred to as “plasticity” (needed for learning). When neurons repeatedly communicate with each other, the connections between them strengthen, making it easier for them to communicate in the future. Conversely, if neurons rarely communicate, their connections weaken and may disappear altogether.

### Major attributes of the human brain that are important for neuromorphic processors:

- A highly interconnected network of neurons
- Communication between neurons only occurs when a threshold event occurs (spike)
- Physical interconnect (synapse) between neurons will either strengthen or weaken depending upon the neuron activity (plasticity) – this is the basis of learning
- Memory is not stored in a single neuron but through a distributed network of connections (synapses)

[4] See for example: https://qbi.uq.edu.au/memory/where-are-memories-stored

# What is a Neuromorphic Computer?

## A neuromorphic computer relies on artificial neurons and synapses to perform computations, **directly mirroring the brain's fundamental building blocks**.

The comparison is imprecise since, at the most basic level, a brain is an electrochemical device, whereas a neuromorphic system relies on semiconductor physics. In addition, the brain is still not well understood, so a neuromorphic system is likely to represent an incomplete (or wrong) model for how the brain processes data.

There are no standards defining how to architect a neuromorphic processor, given the low technical and product maturity of the field. The lack of standard architecture slows progress and market adoption since each instantiation involves unique components and operations. Some developers are representing the features of a neuromorphic system in digital logic, others in mixed-signal (analog and digital) logic, or optical components. Cortical Labs (an Australian startup) is creating a hybrid bio-semiconductor system using actual human neurons embedded on a CMOS chip (see Figure 3). The neurons are alive and can create and modify their synaptic connections with neighboring neurons based upon the underlying electrical signals. To date, they have taught the neuromorphic processor to play the game of Pong[5]. But this is not an easily scalable system. Nor do the neurons currently live long enough to make this a commercially useful product. It is, however, an interesting research tool.

**Figure 3: Cortical labs hybrid human neuron – CMOS chip.**

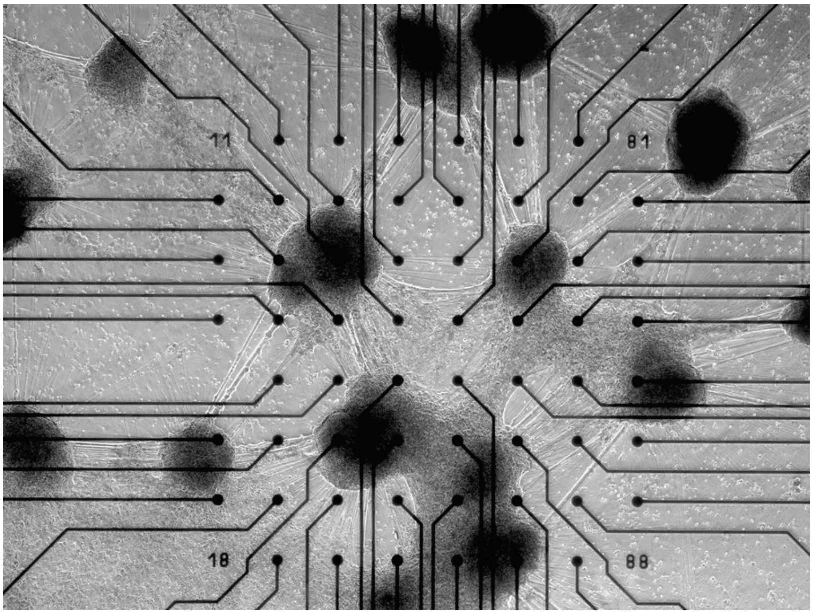


[5] https://corticallabs.com/

**Figure 4: Contrasting a conventional computing architecture (left, referred to as a von Neumann architecture) versus a neuromorphic architecture (right).**

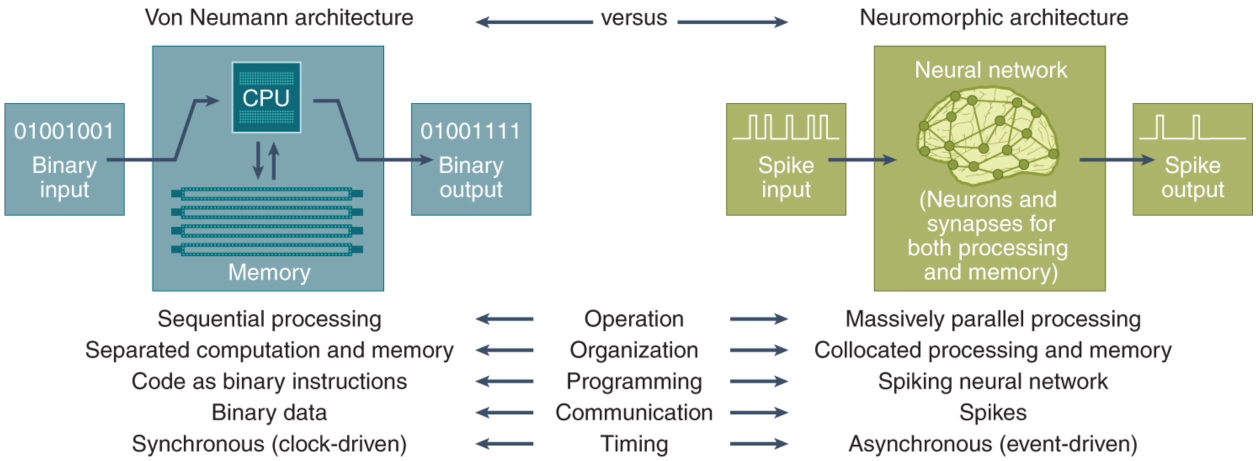


Despite the lack of standard architecture, there exist common attributes across neuromorphic processors, as shown in Figure 4[6]. For example, all neuromorphic systems (for the sake of this paper) will possess these characteristics.

These attributes are contrasted with a von Neumann architecture[7] (standard computer on the left in Figure 4) and form the basis of when **a neuromorphic processor becomes the best choice** for a specific application.

[6] Schuman, C., Kulkarni, S., Parsa, M. et al. 2022, *Opportunities for Neuromorphic Computing Algorithms and Applications,* **Nature Computational Science**, **2**, 10.

[7] https://en.wikipedia.org/wiki/Von_Neumann_architecture

# Advantages of a Neuromorphic Computer

## Neuromorphic processors **offer the following operational advantages** over von Neumann architectures (i.e., classical computing architectures):

- **Low power** – At the core of a neuromorphic processor is the neuron, which communicates only when the inputs reach a threshold and "spikes". For many operations, at any moment in time, most neurons are idle. Calculations are therefore asynchronous ("event driven") rather than synchronous, being controlled by a central clock. This consumes very little power (orders of magnitude less than conventional computing) and aligns neuromorphic processors well to address wearable or embeddable, edge, and mobile device applications (e.g., implants, robotics, IOT, drones, autos).
- **Plasticity** – Neuromorphic systems can dynamically adapt and learn from their environment, improving performance over time without explicit reprogramming. This is particularly valuable for unpredictable real-world scenarios, including smart homes/cities, security systems, and autonomous vehicles.
- **Massively parallel processing** – Neuromorphic systems utilise many simple processing elements ("neurons") operating in parallel, enabling efficient processing of complex tasks simultaneously.
- **Low latency** – Time-critical applications (e.g., autonomous driving) require very low latency, implying local processing (i.e., not utilising the cloud) and not becoming network congested or hitting I/O bandwidth constraints.
- **Avoiding the "memory wall"** – A traditional computer has memory and processing physically separate. This is fundamental to the von Neumann architecture. This architectural decision creates two bottlenecks: 1) The instructions need to be "fetched" from memory before being executed sequentially; and 2) The data is often stored externally and needs to be sent to the processor utilising finite bandwidth links. The memory wall is particularly acute for data-heavy applications such as AI learning and inference. Neuromorphic processors have the data either stored in the synaptic links utilising things like memristors, or use small amounts of on-chip local memory ("in-memory compute"), eliminating the need for external memory.
- **Robustness** – The distributed nature of neuromorphic architectures can lead to higher fault tolerance, as the system can continue to function even if some elements fail. Assuming high connectivity and a system that mimics "plasticity", then new synaptic connections will form to replace dead neurons or broken connections. Robustness is highly desirable for life-critical applications like medical implants.

# Potential Applications

## The above **advantages of neuromorphic computing** are particularly **relevant in near-real-time applications** that require low power and adaptability, including:

- **Autonomous Systems:** Autonomous vehicles, drones, and robots require rapid decision-making based on sensor inputs (e.g., camera, lidar, radar). Neuromorphic processors can analyse the environment in real-time, improving navigation, obstacle avoidance, and overall safety.
- **Smart Cameras:** Neuromorphic cameras can perform real-time image processing for surveillance, traffic management, and crowd monitoring, with lower power consumption and extended operational lifespan.
- **Wearable Devices:** Real-time heart arrhythmia detection, continuous health monitoring, and personalised health tracking can be integrated into the device, providing alerts and eliminating reliance on cloud processing, which is crucial for privacy.
- **Industrial IoT:** Factory sensors powered by neuromorphic chips can analyse vibration or acoustic patterns and other data to detect subtle changes indicative of machine failure. This on-device processing allows for fault detection with significantly lower power consumption, optimising maintenance schedules and reducing downtime.
- **Smart Homes:** Neuromorphic chips can enable adaptive climate control systems that learn user habits and preferences without cloud communication while enhancing privacy protection.
- **Cybersecurity:** Neuromorphic systems can detect subtle anomalies in network patterns or user behaviour that might indicate a cyberattack or breach. Their low latency allows for rapid threat detection.
- **Fraud Detection:** By recognising unusual patterns in transaction data, neuromorphic systems can provide efficient and accurate fraud detection in financial applications.
- **Speech and Image Recognition:** Neuromorphic systems may provide energy-efficient, real-time, on-device speech and image processing, particularly in low signal-to-noise environments.
- **Adaptive Prosthetics:** By analysing muscle or neural signals, neuromorphic systems may enable more intuitive control of prosthetic limbs that can learn and adjust to a user's movement patterns over time.
- **AI Pre- or Post-Processor:** A more speculative application is to create a hybrid neuromorphic – AI system. The concept uses a neuromorphic system to efficiently pre-process AI training data to significantly reduce its size (by orders of magnitude) and allow for a less expensive AI data center. Alternatively, it may be possible to leverage the adaptability of neuromorphic systems to perform reasoning used in inference models to rapidly adapt and (perhaps) approach intelligence mimicry. These are speculative since such a hybrid system has yet to demonstrate "neuromorphic advantage" (although IBM and SpiNNcloud systems are working towards inference demonstrations – see the discussion below).

# Development of Neuromorphic Processors

Figure 5[8] shows the rapid progress made in increasing the number of "neurons" in neuromorphic processors. Compare this chart with Figure 6[9] which depicts the number of neurons in biological organisms (with the largest number of neurons being in African elephants – sorry homo sapiens). With the state-of-the-art system being around 1 billion neurons, we are now approximately matching the number of neurons in a parakeet.

Although progress in creating more neurons is impressive, this is only part of the requirements (a necessary but not sufficient condition). Connectivity is still much less than in a biological brain. Emulating the brain's synaptic connectivity density requires a complex 3-D semiconductor fabrication capability that is beyond the state-of-the-art. Given the important role that synaptic connectivity plays in learning and memory, it is likely that the human level of cognition will not be reproducible in the near or even moderate term (i.e., likely decades away).

**Figure 5: Exponential increase in the number of neurons being simulated in neuromorphic processors.**

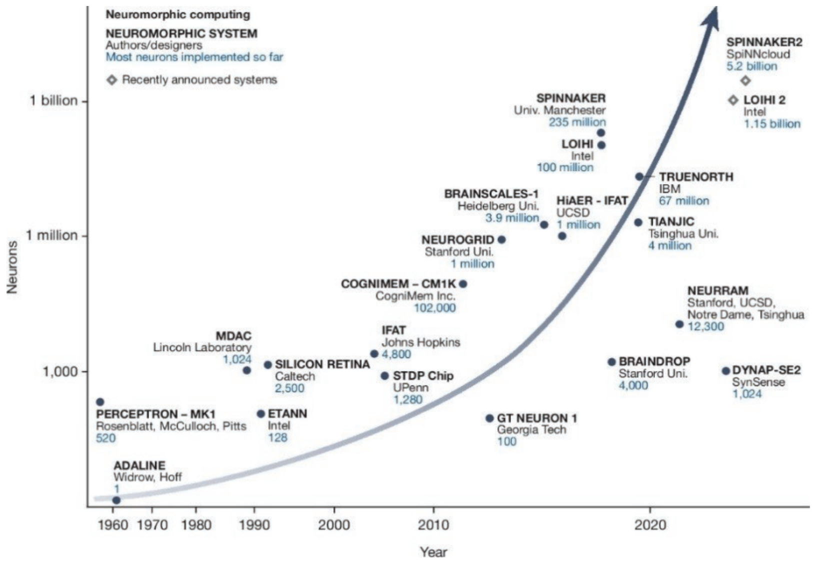


**Figure 6: Typical number of neurons in animals.**

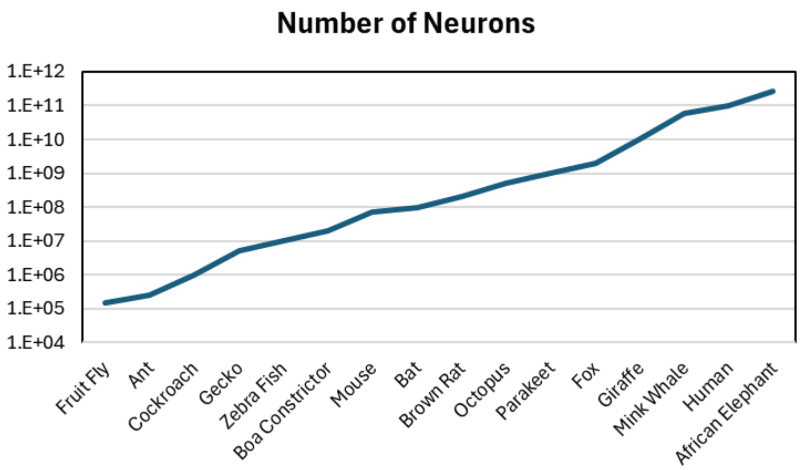


8 Kudithipudi, D., Schuman, C., Vineyard, C.M. et al., 2025, *Neuromorphic Computing at Scale*, **Nature**, **637**, 801.
9 https://en.wikipedia.org/wiki/List_of_animals_by_number_of_neurons

## Despite the technical challenges in scaling to human-level intelligence, **there is optimism around near-term applications in robotics**, autonomous vehicles, and edge computing (e.g., for IOT).

The potential utility of small-scale neuromorphic systems (as well as the frothy AI market in general) is driving investment capital into neuromorphic startups. Crunchbase lists more than 25 neuromorphic startups, with 12 based in the United States, 8 in China, and the rest spread between Europe and Australia[10].

The following table provides a short description of some neuromorphic startups showing a range of hardware for vision, navigation, and edge AI, as well as specialised software that emulates a brain.

| Company | Founded | Product | Headquarters |
|---|---|---|---|
| Opteran | 2020 | S/W mimics insect intelligence for autonomous navigation at the edge | United Kingdom |
| SpiNNcloud Systems | 2020 | Hybrid AI / neuromorphic system. SpiNNaker2 chip (released in 2021) has 152,000 neurons and 152M synapses using 2-5 W. Designed for brain research, robotics, and large-scale hybrid AI models. | Germany |
| Liquid AI | 2023 | Small language models for edge applications using neural networks | USA (spinout of MIT) |
| BrainChip | 2004 | Neuromorphic processor (Spiking Neural Net for Edge AI). The Akida (released in 2022) is a digital chip (~30 mW) for edge AI. | USA (founded in France) |
| Prophesee | 2014 | Neuromorphic vision-based sensor | France |
| SynSense | 2017 | Neuromorphic processor (vision and sensor fusion). The Speck (released in 2022) is a digital chip with an analog visual sensor. Contains 328,000 neurons and runs at ~5 mW. Aimed at computer vision applications. | China (founded in Switzerland) |
| Innatera | 2018 | Analog-mixed signal neuromorphic processor (Spiking Neural Net) for wearables and IOT. The Pulsar (released 2025) uses < 1mW. | Netherlands |

SpiNNcloud Systems is a startup that builds on the Spiking Neural Network Architecture (SpiNNaker) design that came out of the European Union Human Brain Project and managed by the University of Manchester in the UK. At the University of Manchester, they built one of the world's largest neuromorphic systems[11] by combining more than one million neuromorphic processors (see Figure 7). SpiNNcloud has taken this development a step further by creating a hybrid GPU-neuromorphic system with up to 5 million neuromorphic processors that leverage the strengths of both a von Neumann architecture and the energy efficiency of a spiking neural network to create a highly energy-efficient AI inference engine (claimed to be 78x greater energy efficiency than GPUs). Whereas the hybrid architecture is novel, it is unclear if the marketplace will accept the added complexity for the energy savings.

[10] StartUs Insights (https://www.startus-insights.com/innovators-guide/neuromorphic-computing-companies/) lists 119 Neuromorphic Computing Startups from around the world.

[11] https://www.eenewseurope.com/en/spinnaker-neuromorphic-supercomputer-reaches-one-million-cores-2/

## Intel and IBM have also been **exploring neuromorphic processors** for decades.

Intel[12] developed the Loihi 2 processor, which is a spiking neural net chip available for research. The chip consists of 1 million neurons and up to 120 million synapses with a total power of ~1W. The neurons are digital (programmable digital signal processors) but communicate via asynchronous spike messages. Intel combined 1152 Loihi 2 processors into their full-stack Hala Point neuromorphic computer. Hala Point boasts 1.15 billion neurons with 128 billion synapses and uses only 2600 watts.

**Figure 7: Size of the SpiNNaker 1 million core machine. Not quite the size of a parakeet.**

## The system can be programmed using Lava[13], a **neuromorphic programming language** that Intel initially developed and is now available as open source. Intel's desire is for **Lava to become the de facto neuromorphic programming language** regardless of the neuromorphic processor used.

[12] https://www.intel.com/content/www/us/en/research/neuromorphic-computing.html
[13] https://lava-nc.org/

**Figure 8 IBM NorthPole "inference accelerator" performance compared to NVIDIA GPUs on a 3-billion parameter LLM.**

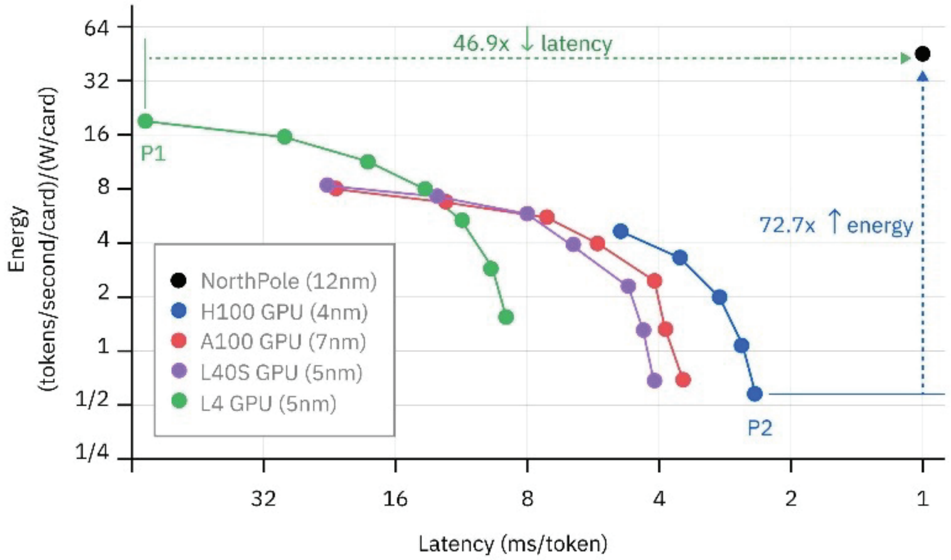


IBM[14] introduced their NorthPole neuromorphic processor in 2023. This chip was optimised to be an AI *inference* machine as opposed to a system that can be generically trained. It mimics some (not all) brain functions[15] including spiking, massively parallel processing, and on-chip memory. The processing precision is only 2, 4, or 8-bit as it tries to replicate low-precision activation signals found within neurons. Referring to NorthPole as an AI inference "accelerator chip", IBM networked 16 NorthPole processors together to perform inference on a 3-billion parameter large language model. In comparison with GPUs, NorthPole had 46.9x lower latency and 72.7x lower energy consumption. Figure 8[16] is from IBM and shows the data for NorthPole compared to NVIDIA GPUs.

Large neuromorphic systems (e.g., SpiNNcloud or Hala Point) are useful for testing new neuromorphic architectures and modeling biological neural systems. However, their size and complexity hide the distinct advantages of a neuro-inspired system. It is unclear whether neuromorphic systems can compete successfully as an AI training pre-processor or AI inference accelerator against classical computing (von Neumann architecture), given the investments in entrenched architecture. This needs to be watched since inference, in particular, has become the dominant AI use (and the largest CapEx outlay). If the efforts of SpiNNcloud or IBM gain traction, we will want to be fully cognisant of the potential market size and speed of adoption.

[14] https://research.ibm.com/blog/northpole-ibm-ai-chip
[15] https://open-neuromorphic.org/blog/northpole-ibm-neuromorphic-ai-hardware/
[16] https://research.ibm.com/blog/northpole-llm-inference-results

# Ecosystem Impediments for the Adoption of a Neuromorphic Computer

Despite the advantages of neuromorphic computing, particularly for edge AI applications, there is **currently little market adoption**. In addition to the technical immaturity of the neuromorphic systems, **the supporting ecosystem is still forming**.

For example, standards do not exist for interfacing with neuromorphic processors (either with classical computing elements or between neuromorphic chips). Software tools and libraries for setting the initial synaptic weights are mostly unique to the specific architecture, and open-source development tools like Lava have yet to become widely adopted.

Foundries have yet to put neuromorphic elements into their standard process design kits (PDKs), making existing systems no more than "research prototypes" and likely with low yield. Memristors (for example) are resistor-like elements that maintain memory of the charge that flowed through them. They are not standard computing elements but can mimic the behavior of neurons and synapses, making them important elements for some neuromorphic architectures.

Human capital is sparse with few experts trained in developing algorithms or using neuromorphic systems. Although the number of academic research papers is increasing (dare I say "spiking"), the commercial expertise is still fairly limited, with only a small cadre of sufficiently trained practitioners. As the market use cases become established, the human capital should develop – but it is currently lacking.

Similar to AI, neuromorphic systems will produce an output, but without an explanation. Trust in the response is lacking an auditable logic trail. This is particularly critical for real-time decisions that may create legal liability (e.g., medical embeddable devices or autonomous operation). Understanding the decision logic of a neuromorphic processor is still an early-stage research problem, and until trust in the system is established, adoption will have significant headwinds.

The impediments listed above are typical for an early technology development. Communities of interest will form and address standards and software development tools. If commercial markets are identified, then foundries will create PDKs with neuromorphic elements. And where money is to be made, human capital will follow. The toughest impediment will likely be explainability, but if AI is a good analog, progress will be made and likely will achieve a level of explainability sufficient for many (perhaps not all) applications.

# Adoption Timeline

## **Neuromorphic computing technology is maturing rapidly.** In just the past 30 years, the number of neurons in a neuromorphic system has increased by >1,000,000-fold.

Multiple architectures are being assessed, and this field continues to be a growing focus in academia. The major impediment to market adoption is not the ecosystem deficiencies discussed above, but the lack of a proven “killer application”.

Assuming neuromorphic systems are successfully productised, two likely scenarios will drive market adoption timing:

1. Killer app (assumed in AI data centres) which will create a large initial demand and fast manufacturing maturation. Inference for image applications (e.g., videos) is a likely candidate to demonstrate early neuromorphic advantage – although this is yet to be demonstrated. As the cost and complexity of designing and operating neuromorphic systems fall, they will then be broadly adopted. If AI data centres drive the demand, it is still likely to be 5+ years to prove out the investment thesis prior to adoption. A rapid demand increase is then possible in the 5-10 year timeframe. Following the AI data centre adoption, wearables are likely the next adopter due to the low risk, followed by IOT, robotics, implantable medical devices, and then autonomous operations.
2. Niche applications leveraging the low power and plasticity in neuromorphic processors are an alternative adoption timeline. IOT and wearables may be the first adopters if performance improvement warrants the development cost. Otherwise, military autonomous systems may drive adoption, creating a sufficiently sized market that lowers the cost for non-military applications. In this scenario, widespread adoption will likely be more than 10 years away.

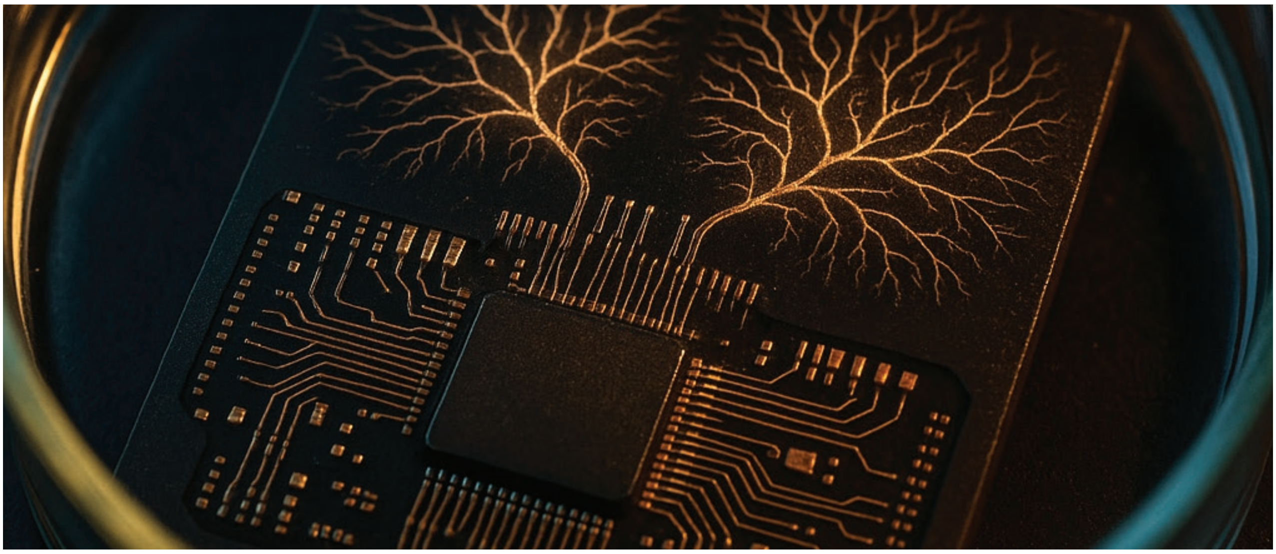

## ICM HPQC Fund

The ICMGF VCC – ICM HPQC Fund (ICM HPQC) invests in the next generation of computing including high-performance computing, quantum technology and new data-centre technologies. ICM HPQC is a sub-fund of ICMGF VCC, a Variable Capital Company (VCC) managed by ICM Global Funds Pte. Ltd (ICM Global Funds).

Capitalising on the ICM Group's experience in novel technologies like AI and quantum computing, as well as its long history of telecommunications and data centre infrastructure investing, ICM HPQC invests in companies that bridge the gap between the demand for more powerful, energy efficient, and secure computing and the ability of existing systems to deliver it.

## ICM Global Funds

ICM Global Funds is a Singapore-based, licensed fund management company dedicated to serving accredited and institutional investors, with products spanning a diverse range of asset classes, including private debt, convertibles, and both private and public equity instruments.

ICM Global Funds is a wholly-owned subsidiary of ICM Limited (ICM), headquartered in Bermuda. Regulated by the Monetary Authority of Singapore, ICM Global Funds holds a Capital Markets Services license. It primarily manages funds through an umbrella Variable Capital Company (VCC) structure which offers clients flexibility, tax efficiency, and privacy.

## About ICM Limited

ICM Limited (ICM) is an entrepreneurial asset management group. Its expertise enables it to offer an innovative, insightful approach beyond traditional investment strategies.

ICM's origins date back to 1988, and it has evolved to become a global organisation, combining the established values and independence of a family office, the measured critique and high engagement of a private equity house, and the patient vision and loyal support of a founding investor. Alongside these qualities, it is increasingly recognised for finding right-sized opportunities in specialist sectors where it can have a meaningful role as an investor. ICM's sectors of focus include: Digital Assets, Financial Services, Global Corporate Bonds, Mobility, Resources, Technology, and Utilities & Infrastructure.

Visit its website **www.icm.limited**

Contact at **info@icm.limited**

## Disclaimer

The document has been prepared by ICM Global Funds Pte Ltd ("ICMGF") for information only. This document should not be relied upon as financial advice or constitute a recommendation and should not be considered an offer or solicitation to deal in ICM HPQC Fund (Registration number T22VC0112B-SF003) (the "Sub-fund"). Accordingly, this document does not constitute or form part of, and should not be construed as, an offer, invitation, or inducement to purchase or subscribe for any securities, nor shall it or any part of it form the basis of, or be relied upon in connection with, any contract or commitment whatsoever in any jurisdiction. Any discrepancies between this document and the Private Placement Memorandum ("PPM"), the terms and conditions outlined in the PPM shall prevail.

The document does not consider the specific investment objective, financial situation, and particular needs of any specific person. You should before making any investment decision, consider the merits or suitability of the fund, the relevant sub-fund or its securities described in this presentation to your particular situation, which should be independently determined by yourself. Any such determination should involve, among other things, an assessment of the legal, tax, accounting, regulatory, financial, credit, and other related aspects of the securities. If you have any doubt about the investment, you should consult an authorised person specialised in advising on investments.

This document is confidential and is being distributed to selected recipients only. It may not be reproduced (in whole or in part), distributed, or transmitted to any other person without the prior written consent of ICMGF. No representation or warranty, express or implied, is made as to, and no reliance should be placed on, the fairness, accuracy, completeness, or correctness of the information, or opinions contained herein. Neither ICMGF, nor any of ICMGF's representatives, including, if relevant, any of ICMGF's affiliates or its directors, officers, agents, etc., shall have any responsibility or liability whatsoever (for negligence or otherwise) for any loss howsoever arising from any use of this document or its contents or otherwise arising in connection with this document. Neither ICMGF nor any other person is under an obligation to keep current the information contained in this document.

This fund will primarily invest in unlisted early-stage companies, which typically have a higher risk of default and illiquidity than later-stage companies and listed equities. Past performance is no guide to future returns. The value of investments and the income from them may go down as well as up, and investors may not get back the full amount they originally invested.

This document has not been approved by the Monetary Authority of Singapore. This document is intended for Accredited Investors or Institutional Investors only, as defined in the Securities and Futures Act 2001.

ICM HPQC Fund is a registered Sub-fund of the ICMGF VCC (the VCC), a variable capital company incorporated in the Republic of Singapore. The assets and liabilities of ICM HPQC Fund are segregated from other Sub-funds of the VCC, in accordance with Section 29 of the VCC Act.