Mapping the Molecular Brain: A New Science of Computation

Two Sentence Summary: We will decode the brain's computational architecture at the molecular level to create the first true blueprints for building radically more energy-efficient computational hardware and continuously learning AI systems. We aim to engineer a new science of computation that accelerates scientific discovery, enables cures for brain disorders by scaling up brain mapping, and enhances the economic momentum of AI and computation.

The big 4 Heilmeier questions

What are we trying to achieve? We will decode the brain's computational architecture at the molecular level to finally translate the brain's existence proof into engineering blueprints that make it possible to build orders of magnitude more efficient computational systems.

How is it done today? What are the limits of current practice? Current brain mapping technology can only access broad brain region activity (e.g. MRI) or brain circuit diagrams (e.g. electron microscopy). These approaches do not reveal what function each connection in a brain circuit performs — we are missing the brain's molecular computational blueprint.

What has changed? What is new in our approach and why will it succeed? In the past 1–2 years, molecular brain mapping technology based on light microscopy has emerged that allows us to characterize brain circuits at the enormous information depth (dozens of information channels at synaptic resolution) required to create a blueprint of their architecture.

Who cares? If successful, what difference will it make? A virtuous cycle of building orders of magnitude more efficient computers from brain blueprints allows us to decode the brain's computational capabilities using AI systems capable of recursive learning. This would accelerate scientific discovery and create economic benefits from orders of magnitude cheaper AIs to finally curing the first brain disorders (global economic cost per year: >$5 trillion).

The other 4 Heilmeier questions

What are the risks and bottlenecks? The two major risks are: (1) a lack of molecular brain mapping creators due to gaps in deployed technology; (2) a lack of accessible infrastructure to support this instrumentation-intensive research. Both risks can be addressed by creating a BBN/MOSIS-style platform for developing and performing molecular brain mapping.

How much will it cost? £16.8M CapEx to equip creators with the necessary technology base (in a platform model, for cost-saving); £3.1M OpEx per year over five years to perform the work.

How long will it take? One year for "last mile" tech dev and charting the course; two years to map brain circuits molecularly; two years to translate the findings and build *in silico* blueprints.

What are the mid-term and final exams to check for success? The exams by technical area (TA) are: (TA1) deployable imaging and data analysis pipelines; (TA2) a list of molecular factors we need to measure to understand how brains compute; (TA3) molecular brain maps of complex biological circuits and derived molecular blueprints for *in silico* architectures; (TA4) *in silico* chips built using the molecular blueprints obtained.

Translating the brain’s existence proof into better computation

The brain is existence proof for orders of magnitude more efficient computation

Moore’s Law — transistors per integrated circuit doubling every two years — has no longer held since the early 2000s. Transistor sizes are constrained by physical limitations in atomic scale and quantum effects. **To continue scaling capabilities, we need a new science of computation.** We do not have to invent this new science — **the brain is existence proof that more efficient computation is possible**: (1) a human brain, from birth to adulthood, consumes just 0.0002% of the energy it took to train the GPT-4 model; (2) whereas AI models require millions of images to tell a dog from a cat, human children learn the distinction heuristically from encountering just a few instances; (3) while AI models are poor learners after training and models collapse during recursive training, human brains learn continuously and recursively.

“Wiring diagram” brain maps are insufficient to decode how the brain computes

We have been able to **map functional brain areas** since the mid-1800s and continue to do so via technologies such as **MRI**, and we have been able to **map the synaptic wiring diagrams of brain circuits** since the 1980s via **electron microscopy**. While **neuromorphic chips** such as IBM’s NorthPole or Intel’s Loihi 2 demonstrate promise in niche applications, we have not succeeded in translating our current understanding of the brain into the kinds of computational advances we know are possible from the brain’s existence proof. We are clearly missing something — **molecular information is inaccessible** in electron microscopy and MRI imaging.

Molecular brain mapping allows us to create blueprints of the brains’ computational architecture

In the past 1–2 years, **using light microscopy to obtain neural wiring diagrams on the level of electron microscopy** has become possible based on a new technology called **expansion microscopy**. Light microscopy unlocks access to the fluorescence-based **toolboxes of molecular biology**, allowing us **to add molecular information to wiring diagrams** — from the localization of synaptic proteins to gene expression. We have obtained all other layers of readouts that are conceivable in brain mapping, from functional readouts of brain areas to synaptic wiring diagrams, and the molecular layer is the only one still missing — it thus seems extremely promising that we can use this technology to finally succeed in **creating molecular blueprints of brains that translate to improved *in silico* architectures.**

A virtuous cycle of brain mapping and computational engineering for human benefit

Once we manage to translate molecular brain map circuit blueprints into radically new *in silico* computational architectures, this new science of computation will create numerous benefits. (1) **Orders of magnitude more energy efficient computation** will allow us to scale AI and miniaturize AI-powered devices into wearables and embodied robotics. (2) These more powerful computational systems will allow us to **develop AIs that learn and adapt** and rely less on human intervention. For example, mapping a single neuron currently requires 30 hours of manual human intervention, at a cost of £1,000 per neuron — and there are 100 billion neurons in the human brain. (3) Accelerating our ability to map the brain via computational architectures based on the brain will thus create a **virtuous cycle of discovery and engineering** that may finally allow us to **cure the first brain disorder — a feat never achieved** — and alleviate their crushing human and economic cost — estimated at >$5 trillion per year globally.

Accelerating the field of molecular brain mapping

Technical Area 1: building a platform for discovery

There are **two bottlenecks** to creating computational blueprints from molecular brain maps. (1) We need **“last mile” tech development to scale** core technologies enabling molecular connectomics — expansion microscopy, circuit mapping, and molecular multiplexing. (2) We need to **develop and deploy specialised imaging and data handling technology**, as commercially available technology is ill-suited to acquire and handle molecular brain maps.

The first bottleneck can be addressed by continuously performing design/development–test loops to **translate the latest scientific progress into deployable technology.** A historic model for this is the firm **Bolt Beranek & Newman (BBN)**, an MIT spin-out that enabled DARPAnet in the 1970s. The second bottleneck can be addressed by lowering the cost of brain mapping by **outsourcing the CapEx-intensive technology base from academic creators** and providing **brain mapping as a service.** A historic model for this is **MOSIS**, a DARPA programme operating since the 1980s to lower access hurdles to custom *in silico* circuit prints.

Both of these approaches can be unified by setting up a (centralized or de-centralized) **platform for molecular brain mapping** that **generates and translates new technology**, and provides access to that technology to the scientific and industrial community to produce brain maps. An initial investment of £32.3M over 5 years (£16.8M CapEx) would be sufficient to set up this platform. Operating costs of £3.1M per year could be covered through (1) government and philanthropic grants, (2) industry user fees, and (3) licensing of imaging and data processing technology invented by platform creators. **This financing outline is modeled on how BBN and MOSIS succeeded in perpetuation.**

Technical Area 2: molecular brain mapping proof-of-concept, modelling, and exploration

To accelerate uptake of this new field, we will set out to prove that we can obtain a fundamentally new **understanding of brain computation from molecular brain mapping.** To this end, we will tackle two biological problem complexes, selected for (1) their small scope, enabling fast results from a focused effort, within 2 years or less, and (2) their powerful scientific breakthrough potential, generating a watershed moment for brain mapping.

Molecular model of the 8-neuron dopamine network governing C. elegans *locomotion & learning*

Dopamine is a rare modulatory neurotransmitter with an outsized influence over behaviour — both in humans and in *C. elegans*. The entire *C. elegans* dopaminergic system — controlling movement and learning — is only 8 neurons, all converging in the small pharyngeal nerve ring. Mapping this small circuit with high behavioural significance that can be modulated via training, genetic means and pharmacological intervention will allow us to model the whole circuit computationally with variable parameters to predict and **link circuit and behaviour.**

Drosophila *larva comparative brain atlas to predict behaviour from molecular ground truth*

Despite the entire brain connectome of the *Drosophila* larva being known, **it has remained impossible to predict the function of each neuron from its connectivity matrix alone.** The function of these neurons can, however, be determined physiologically. This makes it an ideal model system for testing the hypothesis that **molecular annotation could constrain a model of a whole brain sufficiently to predict the function of every neuron.** This is an essential waystone towards enabling the goals of phase two of the programme, i.e. generating architectural blueprints that explain all computation performed within a circuit.

Technical Area 3: molecular blueprints of brain computational architectures

Informed by the results of TA2, **we will create the first blueprints of molecular computation that provide the basis for building better *in silico* computational systems.** We will accomplish this by mapping progressively more complex — and more computationally capable — biological circuits, each of which will be linked to a behavioural or other output measure to **connect computational performance and molecular network architecture.**

Map CL1, the first programmable brain organoid computer

The main challenge in creating a circuit blueprint is to **link the mapped molecular network architecture to the network's output** or performance. The **CL1 organoid engine by Cortical Labs is the first programmable "wetware" computer**, i.e. a computer in which living neurons execute user-defined algorithms. This allows us to **benchmark the network's task performance**, and **obtain the molecularly annotated circuit map of CL1.** This approach allows us to assess which molecular targets we need to read in a biological network to link its architecture to its computational performance.

Map the hippocampal trisynaptic circuit to obtain the first molecular blueprint of a brain circuit

Using the molecular target data gained from CL1, we will **assess complex brain circuits in mammalian systems to fully understand biological computation.** A highly conserved and intensely studied circuit is the trisynaptic circuit of the hippocampus. It is simpler — and evolutionarily older — than many cortical circuit motifs, but it is essential to learning and memory. **Understanding learning and memory in a biological brain is critical** for (1) building more efficient *in silico* computational architectures and (2) devising more flexible AI models. (1) **Integrating storage and processing of information is the most likely strategy to produce immediate improvements in computational efficiency** over existing architectures. (2) **Improving post-training learning and information retention in "world models" is the next significant challenge to significantly improve the capabilities of AI models.**

Map a mouse cortical column to obtain the first logic-gate level blueprint of a complex circuit

The cortical column is a **conserved circuit motif of the mammalian neocortex that has been proposed as the brain's primary "circuit unit" of computation**, with more complex brains — such as the human brain — simply packing more cortical columns to improve computational throughput. While the whole picture is certainly more complex than that, the cortical column still stands out as the most promising circuit to map for **understanding complex, parallelizable and massively adaptive biological computation.** We will select a cortical column with known inputs and outputs that is amenable to behavioural assays — such as the mouse barrel cortex, where one whisker maps onto one cortical column via the thalamus, involved in object recognition and spatial learning. This allows us, again, to **link network architecture and performance** — the basis to **construct a blueprint to build an analogous circuit *in silico.***

Technical Area 4: translating blueprints into building better computational systems

To translate molecular brain maps into better computational architectures, we must address additional challenges such as **engineering *in silico* hardware with more neuron-like properties** — such as memristors and analogue activation threshold transistors. From the earliest days of the programme, we will **engage creators and partners in the neuromorphic compute space** to ensure that the molecular parameters we measure are valuable to guiding transformative *in silico* circuit architectures. Active neuromorphic creator recruitment will be timed based on their early feedback and dataset progress.