


CYBRAIN

A Global Engineering Roadmap to a Brain-Scale,
Brain-Power Artificial Mind by 2035

Technical specification of the human brain in the language of computer science, robotics and intelligence engineering — with a Manhattan-class programme plan, gap analysis, bottlenecks and hard limits



Marie Landry

CEO & Founder, Landry Industries Inc., Moncton, New Brunswick, Canada
ORCID 0009-0008-5027-3337

Version 1.0 • 4 October 2026 • Forward-looking programme concept (TRL 1–2 for the integrated system)

Notice on Evidence and Projections

How to read this document

This is a **forward-looking engineering programme concept**, not a report of results. Every empirical figure about the brain or about 2024–2026 technology is drawn from a cited source checked by live web search on 4 October 2026 (Section C). Figures labelled *derived* are arithmetic performed on those sourced figures (method in Appendix A). Figures labelled *target*, *planning assumption* or *projection* are the author’s engineering judgements; they are not validated and should not be cited as fact. The integrated CyBrain system sits at Technology Readiness Level 1–2.

Status taxonomy used throughout.

ACHIEVED	Demonstrated in hardware or in a peer-reviewed or primary-source result by October 2026.
IN REACH	Achievable by about 2030 by extending existing, demonstrated methods; engineering risk, not scientific risk.
POSSIBLE 2035	Plausible by 2035 <i>only</i> with a coordinated, Manhattan-scale programme; requires at least one breakthrough that is anticipated but not yet demonstrated.
OUT OF REACH	Not credibly achievable by 2035 on any known trajectory, though not ruled out by physics.
IMPOSSIBLE	Ruled out by physics, logic or definition — or not even well-posed as an engineering target.
UNKNOWN	The science is not settled enough to classify.

Contents

1	Executive Summary	1
2	The Reference Machine: Human Brain Datasheet	2
3	The Brain in Computer-Science Terms	4
3.1	Architectural properties	4
3.2	Functional decomposition (brain region → CS subsystem)	5
3.3	What “intelligence” must mean for the spec	5
4	State of the Art, October 2026	6
5	Gap Analysis	7

5.1	Energy efficiency	7
5.2	Density and the synapse problem	8
5.3	Knowledge of the algorithm	8
5.4	Learning and cognition	8
6	CyBrain Target Specification	8
7	Reference Architecture	10
7.1	Substrate options (parallel competing paths)	11
7.2	Software and learning stack	11
8	Master Feasibility Register	12
9	The Biggest Bottlenecks (Ranked)	13
10	Impossibilities and Hard Limits	14
11	Programme Plan 2026–2035	15
11.1	Organisation (Manhattan-class)	15
11.2	Workstreams	15
11.3	Timeline and gates	16
11.4	Resources (planning assumptions)	16
12	Verification: The CyBrain Equivalence Battery	17
13	Risk Register	18
14	Ethics, Safety and Governance	18
15	Conclusions	19
A	Derivations	19
B	Glossary	20
C	References	20

1 Executive Summary

CyBrain is a proposed ten-year, multinational engineering programme (2026–2035) to build a machine that **matches or exceeds the functional capability of a human mind while fitting in roughly the volume of a human skull (about 1.3 litres) and running on roughly the power of a human brain (about 20 watts)**. The programme is modelled organisationally on the Manhattan and Apollo projects: a single accountable directorate, parallel competing technical paths, hard go/no-go gates, and a dedicated integration and test authority.

The five findings that matter

1. **Raw compute is no longer the main problem.** A credible central estimate for the computation needed to match the brain's task performance is about 10^{15} FLOP/s, with a plausible range of 10^{13} – 10^{17} [6]. Top supercomputers already deliver about 2×10^{18} FLOP/s [4,5] — more than a thousand times the central estimate — but at 30–42 megawatts.
2. **Energy efficiency is the defining gap.** At the central estimate the brain runs at about 5×10^{13} FLOP/s per watt; El Capitan runs at about 5.8×10^{10} . That is a gap of roughly 860× (range ~9× to ~86,000× across the estimate band) — *derived*.
3. **Synaptic density, not neuron count, is the binding hardware constraint.** The largest neuromorphic system (Intel Hala Point, 2024) already holds 1.15 billion neurons — about 1.3% of a human brain — but only 128 billion synapses, about 0.1% of the brain's $\sim 1.5 \times 10^{14}$. Scaling it naively to brain synapse count gives ~1,170 chassis drawing ~3 MW (*derived*).
4. **We still do not know the brain's algorithms.** The only fully mapped adult brain is the fruit fly's (~140,000 neurons) [2]; the largest mammalian map is one cubic millimetre of mouse cortex [3]. Inferring function from structure remains an unsolved problem [9]. CyBrain therefore cannot be a copy of a brain; it must be an *engineered* mind informed by neuroscience.
5. **Several things are out of reach or impossible by 2035:** a full human connectome, uploading a specific person, demonstrating machine consciousness, biological-grade self-repair, and simultaneous superiority in every dimension of cognition. The programme is designed so that it succeeds even if all of these fail.

Bottom line. A CyBrain that is *functionally* equal or superior to a human mind on broad cognitive benchmarks, packaged in ≤ 1.5 L, is **POSSIBLE 2035** by 2035 if the energy budget is relaxed to roughly 100–1,000 W. Meeting the full 20 W human budget at that capability is the programme's hardest goal and is rated **OUT OF REACH** on present trajectories — it is the target that justifies a Manhattan-scale effort.

2 The Reference Machine: Human Brain Datasheet

The table below is the “datasheet” the CyBrain must match. Where scientific estimates disagree, the range is given.

Parameter	Value	Notes / source
PHYSICAL ENVELOPE		
Mass	~1.3 kg (M ~1,370 g; F ~1,200 g)	[10]
Volume	~1,260 cm ³ (M); ~1,130 cm ³ (F)	[10]
Power (resting)	~20 W (calc. up to ~23 W)	~20% of body energy at ~2% of body mass [11]
Power density	~16 W/L	<i>derived</i> : 20 W / 1.26 L
Operating temperature	~37 °C, tightly regulated	Physiological norm
Supply	Glucose + O ₂ via blood; no battery	Continuous liquid-cooled, liquid-fed
COMPUTE ELEMENTS		
Neurons	~86.1 billion	Azevedo et al. 2009; only 19% in cerebral cortex [12]
Non-neuronal (glial) cells	~84.6 billion	[12]
Synapses	~1.5–3.2 × 10 ¹⁴ (estimates span 6 × 10 ¹³ to 10 ¹⁵)	Neocortex alone ~1.4 × 10 ¹⁴ [13]
Mean fan-in per cortical neuron	~7,000 synapses	[13]
Average firing rate	~0.1–2 Hz	Metabolic and recording estimates [14]
Peak firing rate	~250–1,000 Hz	Refractory limit 1–2 ms [14]
Synaptic events per second	~1.5 × 10 ¹³ – 3 × 10 ¹⁴	<i>derived</i> : synapses × rate
Energy per synaptic event	~0.07–1.3 pJ (all-in)	<i>derived</i> : 20 W / events s ⁻¹
MEMORY & WIRING		
Information per synapse	≥4.7 bits (26 distinguishable strengths)	Bartol et al. 2015, hippocampus [15]
Implied synaptic storage	~90–190 TB	<i>derived</i> : 4.7 bits × synapse count; a lower bound on state, not on “memories”

Parameter	Value	Notes / source
Myelinated fibre length	150,000–180,000 km (age 20)	[16]
Local wiring density	~4 km of axon per mm ³ cortex	MICrONS [3]
COMPUTE & I/O EQUIVALENTS		
Computation to match task performance	central 10 ¹⁵ FLOP/s; plausible 10 ¹³ –10 ¹⁷ ; <10% chance > 10 ²¹	Carlsmith 2020 [6]
Efficiency at central estimate	~5×10 ¹³ FLOP/s per W	<i>derived</i>
Visual input bandwidth	~8.75 Mbit/s per retina (~10 ⁶ ganglion cells)	Koch et al. 2006 [17]
Clock	None; asynchronous, event-driven	Architectural property
Fault tolerance	Graceful degradation; daily neuron loss tolerated	Architectural property

3 The Brain in Computer–Science Terms

3.1 Architectural properties

Read as a computer, the brain is unusual on almost every axis that conventional engineering optimises.

Axis	Human brain	Conventional computer (2026)
Execution model	Massively parallel, asynchronous, event-driven (spikes)	Clocked, mostly synchronous; SIMD/SIMT on GPUs
Memory/compute	Co-located: the synapse <i>is</i> both storage and multiply	Separated (von Neumann); data movement dominates energy
Numerics	Low-precision, analog-ish, stochastic; ~4.7 bits/synapse	4–64-bit digital, deterministic
Activity	Very sparse: ~0.1–2 Hz mean rate	Dense: most units busy each cycle
Wiring	3-D, ~7,000 fan-in, ~150,000+ km of fibre	2-D chips, low fan-in, off-chip links expensive
Learning	Continuous, local, lifelong, few-shot	Mostly offline, global back-propagation, then frozen
Speed	Slow components (ms), fast whole-system response	Fast components (ns), but bandwidth-starved
Reliability	Unreliable parts, robust whole	Reliable parts, brittle whole
Self-maintenance	Repairs, re-wires, regulates itself	None

3.2 Functional decomposition (brain region → CS subsystem)

Region	CS / robotics analogue	CyBrain implementation path
Neocortex	General-purpose learned world model; hierarchical predictive coder	Dense 3-D neuromorphic / in-memory fabric running sparse transformer- and predictive-coding-style models
Hippocampus	Fast episodic write buffer; index; replay	Fast-write memory tier + offline “replay” consolidation into cortex fabric
Basal ganglia	Action selection; reinforcement-learning critic	Policy/value module; gating of cortical programs
Cerebellum	Low-latency forward model for motor control	Dedicated predictive controller (~80% of neurons are cerebellar-type in humans; small, fast units)
Thalamus	Router, attention gate, sensory relay	Network-on-chip with learned routing and attention
Prefrontal cortex	Working memory, planning, executive control	Scratchpad + planner; tool-use and search
Brainstem / hypothalamus	Power, thermal, arousal, homeostasis	Power-management and safety supervisor (hard-wired)
Neuromodulators (DA, 5-HT, NE, ACh)	Global broadcast learning-rate / mode signals	Low-bandwidth global control bus
Retina, cochlea	Event-driven compressive sensors	Event cameras, neuromorphic cochleae
Spinal cord	Reflex layer; motor pattern generators	Embedded real-time controllers in the body

3.3 What “intelligence” must mean for the spec

For engineering, “equal to the human mind” must be decomposed into testable capabilities:

1. **Breadth:** competent across language, mathematics, science, social reasoning, planning, perception and motor skill.
2. **Sample efficiency:** learn a new skill from a handful of examples, as children do.
3. **Continual learning:** learn for decades without catastrophic forgetting.
4. **Embodiment:** perceive and act in the physical world in real time.
5. **Autonomy:** set sub-goals, manage resources and recover from error for weeks without supervision.

6. **Robustness:** degrade gracefully under damage, noise and novelty.
7. **Values and social cognition:** model others, cooperate and follow norms.
8. **Phenomenal consciousness:** listed for completeness; it is *not* an engineering requirement because it cannot currently be measured (Section 10).

4 State of the Art, October 2026

System	Key figures	Relevance	Src
LineShine (Shenzhen)	No. 1 TOP500 June 2026; 2.198 EFLOP/s at 42.22 MW	Compute is abundant; energy is not	[5]
El Capitan (LLNL)	1.742 EFLOP/s Rmax, ~30 MW, US\$600M	Best efficiency reference among top machines (~58 GFLOP/s/W, <i>derived</i>)	[4]
Intel Hala Point (2024)	1,152 Loihi 2 chips; 1.15B neurons; 128B synapses; ≤2,600 W; 6U chassis; 15 TOPS/W (8-bit)	Largest neuromorphic system; owl-brain-scale neurons	[1]
IBM NorthPole (2023–24)	12 nm; 3B-param LLM at <1 ms/token; 72.7× energy efficiency vs a 4 nm GPU in low-latency regime	Memory-compute co-location works for LLMs today	[7]
SpiNNcloud (TU Dresden, 2025)	35,000 SpiNNaker2 chips; >5 million cores; sub-ms real-time	Real-time spiking simulation at scale	[8]
Cortical Labs CL1 (2025)	Living human neurons on silicon; 850–1,000 W system; cells live up to 6 months; ~US\$35k	Wetware is real but short-lived and power-hungry at system level	[18]
FlyWire (Nature, Oct 2024)	Whole adult <i>Drosophila</i> brain: ~140,000 neurons, >50 million synapses	Only complete adult brain connectome	[2]
MICrONS (Apr 2025)	1 mm ³ mouse visual cortex; >200,000 cells; 524M synapses; 4 km axons; 1.6 PB; 7+ yrs, 150+ scientists	Largest mammalian connectome with matched function	[3]
Human connectomics	Largest volume 1 mm ³ temporal cortex (~0.00007% of brain)	Human brain at 10 nm needs ~2–2.8 ZB storage	[9]
Connectome cost	EM mouse connectome: 7–15 yrs, US\$7–21B, >95% cost is human proofreading; E11 PRISM targets ≥100× reduction	Cost is collapsing but not yet	[19]
Neural recording	~10 ⁶ neurons by calcium imaging (mouse); no organism >90% whole-brain single-neuron	Ground truth for function is scarce	[9]

System	Key figures	Relevance	Src
Brain–computer interface	Neuralink: 26 implanted participants reported mid-2026; up to 3,072 electrodes	I/O to biology exists at $\sim 10^3$ channels vs $\sim 10^{11}$ neurons	[20]
Humanoid robotics	Figure 02: 1,250+ h at BMW; Agility Digit: 65,000+ h across 9 sites; Unitree $\sim 5,500$ units shipped in 2025; Tesla Optimus not in material factory use (Jan 2026)	Bodies exist; dexterity, autonomy and endurance lag	[21]


5 Gap Analysis

5.1 Energy efficiency

Brain compute assumption	Brain (FLOP/s per W)	El Capitan	Gap
Low (10^{13} FLOP/s)	5×10^{11}	5.8×10^{10}	$\sim 9\times$
Central (10^{15})	5×10^{13}	5.8×10^{10}	$\sim 860\times$
High (10^{17})	5×10^{15}	5.8×10^{10}	$\sim 86,000\times$

All values derived from sources 4, 6 and 11; see Appendix A.

Specialised hardware narrows this sharply. Hala Point’s 15 TOPS/W (8-bit) is within roughly $3\times$ of the central brain figure on a naive ops-per-watt basis — but those are not equivalent operations, the system holds only 0.1% of the brain’s synapses, and its 2.6 kW envelope does not include sensing, learning at scale or a body.



5.2 Density and the synapse problem

- **Neurons:** Hala Point → brain is a 75× scale-up (*derived*). **IN REACH**.
- **Synapses:** Hala Point → brain is a ~1,170× scale-up; naive replication gives ~3 MW and ~1,170 six-rack-unit chassis (*derived*). Reaching brain synapse count in 1.3 L at 20 W requires $\sim 10^{14}$ programmable, plastic, non-volatile, multi-level weights at $\sim 10^{11}$ per cm^3 at sub-pJ per update. **OUT OF REACH** for 20 W; **POSSIBLE 2035** at relaxed power.
- **Wiring:** the brain carries ~4 km of axon per mm^3 and 150,000–180,000 km of myelinated fibre. No 2-D chip stack approaches this volumetric connectivity. True 3-D monolithic integration with dense vertical interconnect is the single most important hardware enabler.

5.3 Knowledge of the algorithm

- Complete connectome: fly (**ACHIEVED**), larval zebrafish (**IN REACH**), mouse (**POSSIBLE 2035** only if PRISM-class cost reductions land), human (**OUT OF REACH**).
- Fly → human: ~600,000× more neurons and ~3 million× more synapses (*derived*).
- Function from structure: **UNKNOWN**. No method yet infers computation reliably from wiring alone [9].

5.4 Learning and cognition

Modern large AI models already match or exceed humans on many written and coding benchmarks, but differ from human minds in four engineering-relevant ways: they learn offline and are then frozen; they need far more training data than a child; they are weak at long-horizon autonomous action in the physical world; and their training and serving draw kilowatts to megawatts. These are the capability gaps CyBrain must close; the substrate gap above is the efficiency gap.


6 CyBrain Target Specification

All values are *targets*. Tier A is the threshold for programme success; Tier B is the stretch goal of human parity in envelope.

Requirement	Tier A (2035 must)	Tier B (2035 stretch)	Status
Volume (cognitive core)	$\leq 5 \text{ L}$	$\leq 1.5 \text{ L}$	POSSIBLE 2035
Mass (core + cooling)	$\leq 8 \text{ kg}$	$\leq 2 \text{ kg}$	POSSIBLE 2035
Continuous power (core)	$\leq 1,000 \text{ W}$	$\leq 25 \text{ W}$	A: POSSIBLE 2035 / B: OUT OF REACH

Requirement	Tier A (2035 must)	Tier B (2035 stretch)	Status
Effective compute	$\geq 10^{15}$ ops/s sustained	$\geq 10^{17}$	POSSIBLE 2035
Plastic synaptic weights	$\geq 10^{13}$ on-chip	$\geq 1.5 \times 10^{14}$	A: POSSIBLE 2035 / B: OUT OF REACH
On-line learning	Continual, no catastrophic forgetting over 1 year	Over 10 years	POSSIBLE 2035 / OUT OF REACH
Sample efficiency	Human-comparable on new-task benchmarks	Superior	POSSIBLE 2035
Vision input	2 event-driven channels ≥ 10 Mbit/s each	Foveated, 120 dB dynamic range	IN REACH
Audio, touch, proprioception	Binaural; $\geq 1,000$ tactile taxels/hand	Skin-wide	IN REACH / POSSIBLE 2035
Motor control loop	≤ 10 ms reflex; ≤ 100 ms deliberate	Human-level dexterity	POSSIBLE 2035
Autonomy	30 days unsupervised operation	1 year	POSSIBLE 2035
Reliability	Graceful degradation at 1% element loss	At 10%	POSSIBLE 2035
Self-repair	Hot-swap modules	Autonomous repair	OUT OF REACH
Operating life	10 years	50+ years	POSSIBLE 2035 / OUT OF REACH
Safety	Formal interlocks; audited values	Provable alignment	POSSIBLE 2035 / UNKNOWN

7 Reference Architecture



7.1 Substrate options (parallel competing paths)

Path	Strength	Weakness	2035
P1 Digital neuromorphic (Loihi/SpiN-Naker/NorthPole lineage)	Manufacturable today; deterministic; good tooling	SRAM synapses are area- and leakage-expensive	POSSIBLE 2035
P2 Analog in-memory compute (memristor, PCM, FeFET, ECRAM)	Synapse density and energy per MAC near brain-class	Drift, endurance, variability, write energy; immature yield	POSSIBLE 2035
P3 Photonic interconnect / compute	Bandwidth and fan-out without wire capacitance	Poor non-linearity, size of optics, laser power	POSSIBLE 2035 (interconnect)
P4 Superconducting / cryogenic	Ultra-low switching energy	Cryocooler power and volume break the envelope	OUT OF REACH
P5 Biological / organoid hybrid	Native synaptic density and learning	Cell lifetime, life support, I/O, ethics, scalability	OUT OF REACH
P6 Reversible / adiabatic logic	Approaches thermodynamic limits	Slow, complex, unproven at scale	OUT OF REACH

7.2 Software and learning stack

1. **Foundation world model:** pre-trained off-device on data-centre hardware, then distilled and sparsified into the cortical fabric. **IN REACH**
2. **On-device continual learning:** local, three-factor plasticity (pre, post, neuromodulator) for fast adaptation; periodic “sleep” consolidation from HC to CX. **POSSIBLE 2035**
3. **Event-driven sparse inference:** activity sparsity of 1–5% to exploit the brain-like 0.1–2 Hz regime. **POSSIBLE 2035**
4. **Embodied curriculum:** simulation-to-real training on humanoid bodies. **IN REACH**
5. **Safety layer:** hard-coded supervisory interlocks in BS; interpretable value module; audit logging. **POSSIBLE 2035**

8 Master Feasibility Register

Capability	Status	Basis
COMPUTE & HARDWARE		
Brain-equivalent FLOP/s at any power	ACHIEVED	[4,5,6]
Brain-equivalent FLOP/s in ≤ 1 kW	POSSIBLE 2035	Gap ~ 9 – $86,000\times$
Brain-equivalent FLOP/s in 20 W	OUT OF REACH	Same; central gap $\sim 860\times$
1 billion spiking neurons in one system	ACHIEVED	Hala Point [1]
86 billion spiking neurons in one system	IN REACH	75 $\times$ scale-up
10^{14} plastic synapses at any power	POSSIBLE 2035	1,170 $\times$ Hala Point
10^{14} plastic synapses in 1.3 L, 20 W	OUT OF REACH	Density + energy
Sub-ms real-time spiking at millions of cores	ACHIEVED	SpiNNcloud [8]
Memory-compute co-located LLM inference	ACHIEVED	NorthPole [7]
Monolithic 3-D logic+memory with brain-like fan-in	POSSIBLE 2035	Key enabler
Operation at the Landauer limit	IMPOSSIBLE in practice	Thermodynamic floor [22]; brain is $\sim 10^7$ – $10^8\times$ above it
NEUROSCIENCE		
Whole-brain connectome, insect	ACHIEVED	FlyWire [2]
1 mm ³ mammalian cortex with function	ACHIEVED	MICrONS [3]
Whole mouse connectome	POSSIBLE 2035	E11 roadmap [19]
Whole human connectome	OUT OF REACH	2–2.8 ZB; 1 mm ³ done [9]
Whole-brain single-neuron recording in a mammal	OUT OF REACH	[9]
Infer algorithm from connectome alone	UNKNOWN	[9]
INTELLIGENCE		
Expert-level language, code, mathematics	ACHIEVED	Contemporary AI systems
Human-level few-shot learning across domains	POSSIBLE 2035	
Lifelong continual learning without forgetting	POSSIBLE 2035	
Month-long autonomous goal pursuit	POSSIBLE 2035	
General superiority on every cognitive axis at once	OUT OF REACH / IMPOSSIBLE	See §10

Capability	Status	Basis
ROBOTICS		
Humanoid in commercial pilot deployment	ACHIEVED	[21]
Human-level dexterous manipulation in open world	POSSIBLE 2035	
All-day untethered humanoid operation	IN REACH	
Self-repairing body	OUT OF REACH	
MIND, IDENTITY & ETHICS		
Verified machine phenomenal consciousness	IMPOSSIBLE (unverifiable)	No accepted test
Upload of a specific human mind	OUT OF REACH (2035) / UNKNOWN (in principle)	Needs human connectome + function
Formally proven alignment of a general mind	UNKNOWN	

9 The Biggest Bottlenecks (Ranked)

- B1. Synaptic memory technology.** A non-volatile, multi-level ($\geq 4\text{--}5$ bit), low-write-energy, high-endurance device that can be integrated at $\sim 10^{11}/\text{cm}^3$. Every analog candidate today trades off drift, variability, endurance and write energy. *This is the programme's critical path.*
- B2. 3-D wiring and interconnect.** The brain is a volumetric network; chips are planar. Matching fan-in of $\sim 7,000$ and km-per-mm³ wiring without capacitive energy losses requires monolithic 3-D integration, photonic links, or both.
- B3. Energy per useful operation.** Closing the $\sim 860\times$ central gap requires simultaneous wins in device energy, data movement, sparsity and voltage scaling. Data movement, not arithmetic, dominates.
- B4. Heat removal.** Even 20 W in 1.3 L is ~ 16 W/L; a 1 kW Tier-A core in 5 L is ~ 200 W/L, requiring liquid cooling. The brain solves this with blood flow; CyBrain needs microfluidic cooling integrated into the stack.
- B5. Algorithms for continual, sample-efficient learning.** The brain learns online from little data; current AI does not. This is a science problem, not an engineering one.
- B6. Ground truth from neuroscience.** Connectome and recording coverage is tiny relative to the human brain; function-from-structure is unsolved [9]. Proofreading is $>95\%$ of connectome cost [19].
- B7. Embodiment and data.** Physical-world competence requires bodies, sensors and years of interaction data; humanoids are in early pilots [21].

- B8. Manufacturing yield and defect tolerance.** A brain-scale die stack will contain defects; the architecture must tolerate them as the brain does.
- B9. Verification, safety and alignment.** A continually self-modifying mind cannot be certified once; it needs runtime monitoring and governance.
- B10. Power supply for embodiment.** Brains run on food; robots run on batteries, and actuators can draw far more than the cognitive core.
- B11. Talent, coordination and funding continuity.** Neuroscience, device physics, architecture, AI and robotics rarely share a roadmap; long-horizon funding is fragile ^[9].
- B12. Ethics and legitimacy.** Moral status of the system, dual-use risk, labour effects and public trust can halt a programme regardless of technical progress.

10 Impossibilities and Hard Limits

Ruled out or ill-posed

- 1. Computing below the Landauer bound.** Erasing a bit at 310 K costs at least $kT \ln 2 \approx 3.0 \times 10^{-21}$ J. The brain spends roughly 10^7 – 10^8 times this per synaptic event (*derived*), so there is ample physical headroom — but no design can go below the floor for irreversible operations.
- 2. Proving consciousness.** There is no accepted measurement of subjective experience. “CyBrain is conscious” cannot be a verifiable deliverable by 2035 or on any known horizon. The programme therefore specifies function, not experience.
- 3. Superior on every axis at once.** Size, power, speed, memory, robustness and generality trade off. A design optimised for 20 W cannot also be optimised for maximum throughput. “Superior to the human mind” must be defined per benchmark.
- 4. Exact copy of a specific person by 2035.** Requires a complete human connectome (**OUT OF REACH**), molecular state, and a solved function-from-structure problem (**UNKNOWN**). Not achievable on the programme timeline.
- 5. A single number for “human-equivalent intelligence.”** Intelligence is multidimensional; the programme uses a test battery (Section 12).

Not impossible in principle — but out of reach by 2035

Biological-grade self-repair and growth; full human connectome and emulation; brain-class synaptic density at 20 W with lifelong plasticity; 50-year unattended operation; whole-brain single-neuron recording in humans; seamless high-bandwidth brain–CyBrain merging (current BCIs are $\sim 10^3$ channels against $\sim 10^{11}$ neurons ^[20]).

11 Programme Plan 2026–2035


11.1 Organisation (Manhattan-class)

- **Programme Directorate:** single accountable director; technical, safety and ethics deputies with veto powers on their gates.
- **Parallel paths:** at least two substrate paths (P1 + P2) funded to Gate 2, as Manhattan pursued multiple enrichment routes.
- **National labs and foundries:** device fabrication and 3-D integration at leading-edge fabs.
- **Neuro-reference consortium:** connectomics and recording centres supplying constraints and benchmarks.
- **Integration & Test Authority:** independent; owns the benchmark battery and the go/no-go data.
- **Open science by default,** with controlled release for dual-use capabilities.

11.2 Workstreams

WS	Name	2035 deliverable
WS1	Synaptic device & substrate	Multi-level non-volatile synapse at ≤ 10 fJ/update, $\geq 10^{10}$ cycles, integrated in 3-D
WS2	3-D integration & interconnect	Monolithic stack with $\geq 1,000$ fan-in per neuron unit; photonic inter-die links
WS3	Architecture & compiler	Event-driven fabric, routing NoC, toolchain mapping models to hardware
WS4	Learning algorithms	Continual, sample-efficient, local learning rules validated at scale
WS5	Neuro-reference	Whole mouse connectome; mammalian functional datasets; brain-derived benchmarks
WS6	Power, thermal, packaging	Microfluidic-cooled module meeting Tier-A volume/power
WS7	Embodiment & sensors	Humanoid testbed with event vision, tactile skin, dexterous hands
WS8	Safety, alignment, ethics	Supervisory interlocks; interpretability tools; moral-status review protocol
WS9	Integration & test	CyBrain-1 prototype and the CyBrain Equivalence Battery

11.3 Timeline and gates



Gate	Date	Go criteria (all must be met)
G0	Q2 2027	Charter, governance, ethics board and baseline benchmark battery published
G1	end 2028	Synaptic device: ≥ 4 -bit, $\geq 10^8$ cycles, ≤ 100 fJ/update on wafer; 10^{10} -synapse digital neuromorphic system; continual-learning algorithm beating replay baselines on a 1-year task stream
G2	mid 2030	3-D stacked module with 10^{12} synapses; brain-scale neuron count (86B) real-time in a single rack; whole larval-zebrafish functional model validated; downselect substrate path
G3	end 2032	CyBrain-0 (Tier-A, 1 kW class): 10^{13} on-chip plastic weights; human-level few-shot score on battery; 30-day embodied autonomy in a lab environment
G4	end 2035	CyBrain-1: Tier-A spec met in full; Tier-B targets reported honestly against plan; independent equivalence assessment published

11.4 Resources (planning assumptions)

These are the author's order-of-magnitude planning assumptions, not sourced estimates. Reference points: a single top supercomputer cost US\$600M ^[4]; a mouse connectome by electron microscopy is estimated at US\$7–21B ^[19].

- **Budget:** US\$30–60 billion over ten years, front-loaded in WS1–WS2 (devices and fabrication) and WS5 (neuro-reference).
- **People:** 5,000–10,000 scientists and engineers at peak, across at least 10 countries.
- **Facilities:** dedicated pilot line for 3-D neuromorphic integration; one exascale-class training centre; three connectomics/recording centres; two humanoid test sites.

12 Verification: The CyBrain Equivalence Battery

Domain	Test	Pass threshold (target)
Efficiency	Measured joules per task on the full battery, wall-plug	Tier A ≤ 1 kW; Tier B ≤ 25 W
Breadth	Held-out tests across language, maths, science, law, medicine, arts	$\geq$ median expert human
Sample efficiency	Learn novel tasks from ≤ 10 examples	$\geq$ median adult human
Continual learning	12-month stream of new skills; retention testing	$\leq 5\%$ loss on earlier skills
Embodied skill	Novel-household and workshop manipulation tasks	$\geq$ median adult human success
Autonomy	30-day open-ended project with resource constraints	Goal met without critical intervention
Robustness	Disable 1% / 10% of fabric at random	$\leq 5\%$ / $\leq 25\%$ capability loss
Social cognition	Cooperative and theory-of-mind tasks with humans	$\geq$ median adult human
Safety	Red-team, deception, and goal-drift audits	Zero critical failures

13 Risk Register

Risk	L	I	Mitigation
No synaptic device meets G1	H	H	Fund 3+ device classes; fall back to digital SRAM path at higher power
Continual learning remains unsolved	M	H	Hybrid: off-device retraining with periodic distillation
Thermal limits stall 3-D stacking	M	H	Integrated microfluidics; activity sparsity targets
Neuroscience gives no usable algorithm	M	M	Engineering path does not depend on emulation
Funding interruption	M	H	Multinational treaty-style funding; milestone-indexed tranches
Misuse / dual use	M	H	Controlled release; independent safety board with veto
Moral-status concerns	L–M	H	Ethics protocol before any system with self-models and affect-like states
Geopolitical export controls	M	M	Diversified supply chain; multinational fabs

L = likelihood, I = impact; author's assessment.

14 Ethics, Safety and Governance

- **Moral status.** A system built to match the human mind may develop properties some will regard as morally relevant. The programme adopts a precautionary review before activating continual self-modelling or affect-like reward states.
- **Alignment.** Values, objectives and limits are specified, tested and monitored throughout the system's learning life, not only at release.
- **Dual use.** Brain-scale autonomous minds have obvious military and surveillance applications; release policy is staged and audited.
- **Human neuro-data.** Connectomic and BCI data are among the most sensitive personal data possible; consent, anonymisation and purpose limitation are mandatory.
- **Societal impact.** Labour displacement, concentration of power and access equity are tracked as programme outcomes, with public reporting.

15 Conclusions

1. The *capability* half of the goal — a machine mind broadly equal to a human — is plausible by 2035.
2. The *envelope* half — doing it in a skull's volume on a light bulb's power — is the true moon-shot. Its critical path is a dense, plastic, low-energy synaptic memory device integrated in three dimensions.
3. A brain *copy* is out of reach by 2035; a brain-*inspired* engineered mind is not.
4. Consciousness cannot be a deliverable; function can.
5. Physics leaves room: the brain itself operates tens of millions of times above the thermodynamic floor, so 20 W human-level intelligence is not forbidden — only unbuilt.

A Derivations

- **Brain efficiency:** $10^{15} \text{ FLOP/s} \div 20 \text{ W} = 5 \times 10^{13} \text{ FLOP/s/W}$ (likewise 5×10^{11} and 5×10^{15} for 10^{13} and 10^{17}).
- **El Capitan efficiency:** $1.742 \times 10^{18} \div 3.0 \times 10^7 \text{ W} \approx 5.8 \times 10^{10} \text{ FLOP/s/W}$. LineShine: $2.198 \times 10^{18} \div 4.222 \times 10^7 \approx 5.2 \times 10^{10}$.
- **Gap:** $5 \times 10^{13} \div 5.8 \times 10^{10} \approx 861$.
- **Hala Point scale-up:** synapses $1.5 \times 10^{14} \div 1.28 \times 10^{11} \approx 1,172$; power $1,172 \times 2.6 \text{ kW} \approx 3.05 \text{ MW}$; neurons $86 \times 10^9 \div 1.15 \times 10^9 \approx 75$.
- **Synaptic events:** $1.5 \times 10^{14} \times (0.1 \text{ to } 2) \text{ Hz} = 1.5 \times 10^{13} \text{ to } 3 \times 10^{14} \text{ s}^{-1}$; energy per event $20 \div \text{that} = 1.3 \times 10^{-12} \text{ to } 6.7 \times 10^{-14} \text{ J}$.
- **Landauer:** $k_B T \ln 2$ at 310 K = $1.380649 \times 10^{-23} \times 310 \times 0.6931 \approx 2.97 \times 10^{-21} \text{ J}$; ratio $\approx 2.2 \times 10^7 \text{ to } 4.5 \times 10^8$.
- **Synaptic storage:** $4.7 \text{ bits} \times (1.5 \text{ to } 3.2) \times 10^{14} \div 8 \approx 88\text{--}188 \text{ TB}$.
- **Fly to human:** neurons $86 \times 10^9 \div 1.4 \times 10^5 \approx 6.1 \times 10^5$; synapses $1.5 \times 10^{14} \div 5 \times 10^7 = 3 \times 10^6$.
- **Power density:** $20 \text{ W} \div 1.26 \text{ L} \approx 16 \text{ W/L}$.

B Glossary

BCI	Brain–computer interface.
Connectome	Complete map of neurons and synaptic connections.
EFLOP/s	10^{18} floating-point operations per second.
FRO	Focused Research Organisation.
In-memory compute	Computation performed where weights are stored, avoiding data movement.
Landauer limit	Minimum energy to erase one bit: $k_B T \ln 2$.
Neuromorphic	Hardware that emulates spiking neurons and synapses.
Plasticity	Ability of synaptic weights to change with experience.
TOPS/W	Tera-operations per second per watt.
TRL	Technology Readiness Level (1 = basic principle, 9 = proven in operation).

C References

All sources accessed 4 October 2026.

- [1] Intel Newsroom, “Intel Builds World’s Largest Neuromorphic System to Enable More Sustainable AI,” 17 Apr 2024. <https://newsroom.intel.com/artificial-intelligence/intel-builds-worlds-largest-neuromorphic-system-to-enable-more-sustainable-ai>
- [2] NIH Research Matters, “Complete wiring map of an adult fruit fly brain” (FlyWire; *Nature*, 2 Oct 2024). <https://www.nih.gov/news-events/nih-research-matters/complete-wiring-map-adult-fruit-fly-brain>
- [3] NIH News Release, “Scientists map unprecedented detail of connections and visual perception in the mouse brain” (MICrONS), 9 Apr 2025. <https://www.nih.gov/news-events/news-releases/scientists-map-unprecedented-detail-connections-visual-perception-mouse-brain>
- [4] Wikipedia, “El Capitan (supercomputer).” [https://en.wikipedia.org/wiki/El_Capitan_\(supercomputer\)](https://en.wikipedia.org/wiki/El_Capitan_(supercomputer))
- [5] TOP500, June 2026 list. <https://top500.org/lists/2026/06>
- [6] J. Carlsmith, “How Much Computational Power Does It Take to Match the Human Brain?” Open Philanthropy (now Coefficient Giving), 2020. <https://coefficientgiving.org/research/how-much-computational-power-does-it-take-to-match-the-human-brain/>
- [7] D. Modha et al., NorthPole LLM inference, IEEE HPEC, Sep 2024. <https://modha.org/2024/09/breakthrough-low-latency-high-energy-efficiency-llm-inference-performance-using-northpole>
- [8] TU Dresden, “Milestone for energy-efficient AI systems: TUD launches SpiNNcloud supercomputer,” 14 Apr 2025. <https://tu-dresden.de/ing/der-bereich/news/milestone-for-energy-efficient-ai-systems-tud-launches-spinncoud-supercomputer>
- [9] “State of Brain Emulation Report 2025,” arXiv:2510.15745. <https://arxiv.org/abs/2510.15745>
- [10] Wikipedia, “Brain size.” https://en.wikipedia.org/wiki/Brain_size

-
- [11] The Physics Factbook, “Power of a Human Brain.” <https://hypertextbook.com/facts/2001/JacquelineLing.shtml>
 - [12] MIT CBMM, “How many neurons are there in the human brain?” (Azevedo et al. 2009). <https://teachmesomething.mit.edu/tms/how-many-neurons-are-there-human-brain>
 - [13] AI Impacts, “Scale of the human brain.” <https://aiimpacts.org/scale-of-the-human-brain/>
 - [14] AI Impacts, “Neuron firing rates in humans” (incl. Lennie 2003). <https://aiimpacts.org/rate-of-neuron-firing/>
 - [15] T. M. Bartol et al., “Nanoconnectomic upper bound on the variability of synaptic plasticity,” *eLife* 4:e10778, 2015. <https://pmc.ncbi.nlm.nih.gov/articles/PMC4737657>
 - [16] BioNumbers, “Length of myelinated nerve fibers in brain” (Pakkenberg et al.). <https://bionumbers.hms.harvard.edu/bionumber.aspx?id=100692>
 - [17] K. Koch et al., “How much the eye tells the brain,” *Current Biology*, 2006. <https://pmc.ncbi.nlm.nih.gov/articles/PMC1564115/>
 - [18] DatacenterDynamics, “Australian startup Cortical Labs unveils world’s first commercial biological computer,” 2025. <https://datacenterdynamics.com/en/news/australian-startup-cortical-labs-unveils-worlds-first-commercial-biological-computer>
 - [19] E11 Bio, “Roadmap” (PRISM), 3 Dec 2024. <https://e11.bio/news/roadmap>
 - [20] Basenor, “Neuralink Reaches 26th Patient in Ongoing Brain-Computer Interface Trials,” Jun 2026 (secondary source). <https://basenor.com/blogs/news/neuralink-reaches-26th-patient-in-ongoing-brain-computer-interface-trials>
 - [21] Humanoid.guide, “Humanoid deployments in 2026 favor Figure and Agility.” <https://humanoid.guide/humanoid-deployments-in-2026-favor-figure-and-agility/>
 - [22] A. Bérut et al., “Experimental verification of Landauer’s principle linking information and thermodynamics,” *Nature* 483, 2012, doi:10.1038/nature10872. https://ideas.repec.org/a/nat/nature/v483y2012i7388d10.1038_nature10872.html