FUTURE POWER CAPITAL
CAPITAL
BUILDING VALUE · DRIVING INNOVATION · SHAPING TOMORROW

FUTURE POWER CAPITAL — DEEP TECH PROGRAMME

The Long View

A thesis for deep technology, and why a ten-year horizon is an advantage rather than a constraint.

INTRODUCTION

Why this document exists

The Deep Tech Programme page on our website states, in a few paragraphs, what the programme intends to do and who it is built for. This document is the extended version of that same position — the scientific thesis behind it, written for an investor, an industrial partner, or a research founder who wants to understand our reasoning in full before any conversation begins.

Nothing here contradicts or replaces the programme page. Where that page states an intent, this document explains the reasoning behind it in more depth: why we believe hard science deserves a distinct funding posture, what we mean concretely by each scientific domain we intend to look at, how we plan to evaluate a project, and what a ten-year horizon changes about the relationship between capital and a research team.

How to read what follows

Two registers run through this document, and we have kept them visually distinct throughout:

- **Established scientific fact.** Presented in plain body text. These are claims about the state of the art — what has been demonstrated, published, or measured — and we have checked them against public sources rather than writing from memory. Where a timeline or a number is uncertain, we say so.
- **Investor perspective.** Set apart in a shaded box, like the one below. These are our own views on what a given fact means for how capital should behave — not additional facts, and not predictions we are asking anyone to rely on.

> **Investor perspective.** We think this distinction matters more in deep tech than anywhere else in our activity. A scientific claim can be checked against a publication or a measurement; an investment opinion about what that claim is worth cannot. Conflating the two — writing marketing language as if it were data — is, in our view, the single most common way that serious science gets oversold. We would rather be plain about which sentences are ours.

What this document is not

It is not a fund prospectus, an offer of investment, or a list of current holdings. Future Power Capital does not disclose the names of businesses it has invested in, the amounts committed, or the terms of any transaction, as a matter of policy — this document contains none of those things, and none should be inferred from its absence. It does not name any research partner, university, or laboratory we work with, because at the time of writing we do not present any such partnership as established fact. It contains no return figures, no performance projections, and no promise about what any technology described here will be worth. What it does contain is our honest reasoning about a category of investment we believe is structurally underserved, and the method we intend to apply to it.

PART ONE — THE THESIS

Why deep technology, why now

"Deep technology" is a loose label for ventures built on a genuine scientific or engineering advance, rather than on a business model layered over existing technology. The label is not new. What has changed, in the last decade, is the density of foundational advances arriving at roughly the same time, in fields that had each individually spent decades in relative dormancy.

A convergence, not a single trend

Three shifts are documented well enough to state as fact rather than forecast. First, machine learning moved from a narrow toolkit to a general-purpose method for exploring extremely large search spaces — protein conformations, candidate materials, control policies — spaces that were previously intractable to search by trial and error alone. Second, decades of public and private investment in genomics, sequencing, and cell and gene therapy platforms lowered the cost of manipulating biological systems by orders of magnitude compared with the year 2000. Third, the economics of batteries, renewable generation, and computing hardware each crossed cost thresholds during the 2010s and 2020s that made previously theoretical applications commercially arguable for the first time.

None of these three shifts, on its own, would justify a distinct investment category. Together, they change the texture of what a well-run laboratory or engineering team can plausibly attempt, and they do so at a moment when several of the underlying tools — computation, sequencing, materials characterisation — are still getting cheaper every year rather than plateauing.

Why "now" is a specific claim, not a slogan

We are careful with the word "now," because it is the word most abused in technology marketing. Our claim is narrower than it sounds: not that transformative breakthroughs are imminent on any particular calendar, but that the cost of attempting rigorous, well-instrumented research has fallen enough, in several fields at once, that the marginal team with a genuine idea can test it further before running out of money than was true fifteen years ago. That is a statement about the cost of experimentation, not about the certainty of outcomes.

> **Investor perspective.** We do not think every deep-tech field is investable at every point in time, and we say so plainly in the sections that follow: some of the domains we cover are, on our own assessment, further from a fundable inflection point than others. The point of this section is narrower — that the category as a whole deserves a distinct posture from capital, not that every project within it deserves funding.

Why a long horizon is an advantage, not a constraint

Conventional venture capital is built around fund lifecycles of roughly ten years, with most of the capital deployed in the first three to five and pressure to show markable progress well before the underlying science may be ready to support it. That structure is not a flaw in venture capital — it is a rational response to how venture funds themselves are financed, typically by limited partners who need a return within a bounded period. But it means that a fund's incentives and a deep-tech team's actual timeline are frequently mismatched, independent of how good the science is.

Capital that does not answer to a fixed-term limited partnership does not face that mismatch in the same way. It can size a commitment to the actual pace of the work, hold a position through a phase with no obvious commercial marker, and judge progress by the technical milestones a field itself recognises rather than by a fundraising calendar set elsewhere. This is not a claim that patient capital produces better outcomes — that would be a return promise, and we make none. It is a narrower, structural claim: that removing a fixed exit clock removes one source of pressure that has nothing to do with whether the science is working.

PART TWO — THE FINANCING PROBLEM

Why the usual financing cycles do not fit research-intensive companies

Two financing traditions dominate how technical ventures are usually funded: venture capital, and public or philanthropic grant funding. Both were built for a purpose that is not quite the purpose a deep-tech founder actually has, and the mismatch is well documented rather than a matter of opinion.

What venture capital is built to do

Venture capital is optimised to find companies that can grow revenue very quickly once a product exists, because a fund's return depends on a small number of large, relatively fast outcomes within its ten-year life. Research published on European deep-tech financing describes a "missing middle": a stage where the technical risk is still substantial and the capital required to reach industrial scale is already large, sitting between initial research funding and the traction that would justify a typical growth-stage venture round. Multiple independent analyses of deep-tech seed-to-Series-A timelines report that this stage commonly takes around twelve months longer than an equivalent software company, precisely because "traction" for a hard-science company often means a successful pilot or a validated physical result, not a user-growth curve.

What grant funding is built to do

Public and philanthropic grants are generally built to fund research whose primary output is knowledge — a publication, a proof of concept, a dataset — rather than a company. Grant cycles typically run on annual or multi-year application calendars set by the funding body, are usually non-dilutive but also non-recurring in any guaranteed way, and are often awarded to an academic institution rather than to a company that can hold equity, hire commercially, or make binding commercial decisions. A team that has already incorporated and wants to run the fast, iterative, product-shaped work that follows a scientific result often finds that the grant instruments available to it were not designed for that stage.

The gap between the two

	Venture capital	Grant funding	The gap between them
Primary output expected	Revenue growth, market traction	Publications, knowledge, proof of concept	A validated technical result with no market yet

	Venture capital	Grant funding	The gap between them
Typical decision cycle	Weeks to months, deal by deal	Fixed annual or multi-year calls	Continuous technical iteration, no fixed calendar
What is funded	A company with a product	Usually a research institution or PI	An incorporated team without a market-ready product
Tolerance for a negative result	Low — a fund needs winners	High for the specific grant, but funding does not continue automatically	Needs continuity of support through a negative result

> **Investor perspective.** We recognise this gap is not a new observation — it is discussed at length in deep-tech financing literature under labels such as the "missing middle" or the "valley of death." We are not claiming to have discovered it. Our position is narrower: that the gap is real enough, and persistent enough, that a capital structure built specifically to sit inside it — patient, principal, and willing to fund a company rather than a publication — has a distinct role to play, independent of what any individual project inside it eventually achieves.

What independent capital changes, concretely

Because Future Power Capital commits its own capital as principal rather than capital raised from external limited partners with a fixed-term mandate, a Deep Tech Programme commitment does not need to be justified to a committee acting on someone else's behalf, and it does not carry a countdown to an exit set by a fund's own lifecycle rather than by the technology. In practice, that changes three things relative to a conventional venture structure:

- **The size and shape of a commitment can follow the actual cost of the next technical milestone**, rather than being sized to fit a standard round structure built for software companies.
- **A negative or inconclusive technical result does not automatically end the relationship**, if the underlying reasoning for backing the team in the first place still holds — a point we return to directly in Part Five, on accepting scientific failure.
- **Governance can be lighter where the risk is genuinely scientific rather than commercial** — we are not attempting to manage a laboratory's research agenda from the outside, only to fund it responsibly and ask honest questions about progress.

What this does not change

Patient, independent capital is not a substitute for scientific rigour, and it is not immune to the underlying risk that deep-tech research carries. A team can be well-funded on a sensible timeline and still fail to make the science work — that risk is inherent to the category, and no financing structure removes it. What a longer horizon and independent governance can do is make sure that the team's failure, if it comes, is a scientific failure and not a financing-calendar failure.

> **Investor perspective.** We think this is the honest way to describe what patient capital is for. It is not a claim that our approach produces better science, or a higher rate of success — we make no such claim and have no evidence that would support one. It is a claim about removing one specific, well-documented source of pressure — the mismatch between fund lifecycle and research timeline — and being transparent that every other source of risk in deep-tech investing remains fully in place.

PART THREE — THE DOMAINS WE COVER

Five domains, and how we think about each

The Deep Tech Programme is not built around a single technology bet. The five domains below span very different scientific bases, timelines, and failure modes, and we treat each on its own terms. For every domain we set out: the real state of the art as it is documented today, the technical bottlenecks that are publicly known to stand between today's state of the art and a broader commercial or scientific outcome, a plausible maturity horizon where one can be responsibly stated, and the signals we use to separate a credible technical team from one that oversells its position.

> A note on maturity horizons. Where we give a horizon below — for example, "several years" or "toward the end of the decade" — this reflects a synthesis of publicly available technical roadmaps and independent analysis at the time of writing, not a Future Power Capital forecast and never a promise. Public roadmaps for hard technology have slipped before, repeatedly and by years, and we say so explicitly where the historical record shows it.

The five domains

- **Applied artificial intelligence** — beyond general-purpose chatbots, into domain-specific systems for scientific and engineering discovery.
- **Life sciences and biotechnology** — the application of computational methods to biological design, and the manufacturing realities that follow a computational result.
- **Energy and storage** — the technologies competing to succeed today's lithium-ion and combustion-dominated base, including fusion as a longer-horizon case.
- **Advanced materials** — the acceleration of materials discovery by computation, and the much slower process of turning a candidate material into a manufacturable product.
- **Computing and infrastructure** — the physical constraints now binding artificial intelligence progress, and the position of quantum computing on a multi-decade path to practical use.

DOMAIN ONE

Applied artificial intelligence

The public conversation about artificial intelligence is dominated by general-purpose language models. The domain we are actually interested in for the Deep Tech Programme is narrower: the application of machine learning to genuinely hard scientific and engineering search problems, where the model is a tool inside a larger discovery process rather than the product itself.

State of the art

Two facts are well established. First, frontier AI systems have continued to improve on published benchmarks through increased training compute, larger datasets, and — more recently — increased computation applied at inference time to let a model "reason" through a problem step by step. Stanford's 2025 AI Index reports that training compute for frontier models has been doubling roughly every five months, alongside comparable growth in dataset size and energy use. Second, and independently of raw scaling, a growing body of published research — including work from Apple's machine learning research group on the behaviour of large reasoning models — documents that these systems still fail at exact, multi-step computation and reason inconsistently across structurally similar problems, even as they succeed on complex-looking benchmark tasks. Scaling and genuine robustness are not the same axis, and public research now treats that distinction as established rather than speculative.

Where this intersects deep tech specifically

The application that matters most for our purposes is not chat interfaces but AI as a search tool over enormous combinatorial spaces — candidate molecules, candidate materials, control policies for physical systems — where brute-force experimentation was previously too slow or too expensive. This is the same underlying capability driving progress in the biotechnology and materials domains discussed below, which is why we treat applied AI as a horizontal domain rather than a self-contained one.

Technical bottlenecks that are publicly documented

- **Reliability under exact computation.** Published evaluation work shows current reasoning models degrade sharply on problems requiring long, exact, multi-step logic, even when they perform well on pattern-matching-heavy benchmarks — a gap directly relevant to any application where a wrong answer has a real cost.
- **Compute and energy cost.** Training-compute growth on a roughly five-month doubling cycle is not free, and it concentrates capability development among organisations that can fund that growth, which shapes who can credibly compete at the frontier.

- **Distribution shift.** A model trained on existing scientific literature and data can struggle when a genuinely novel discovery falls outside the distribution it was trained on —

precisely the situation deep-tech discovery is meant to produce.

A plausible maturity horizon

Domain-specific scientific AI tools are not a future technology — they are already integrated into materials science and structural biology workflows today, as the next two domains describe. What remains genuinely uncertain, on the public record, is the reliability of general-purpose reasoning systems for autonomous, unsupervised scientific work; independent published evaluations through 2025 do not support treating that as solved.

What distinguishes a credible team from one that oversells

Signal	Credible team	Warning sign
How they describe model limits	States specific failure modes and error rates against a defined benchmark	Describes the model only through its successes, with no discussion of failure cases
Validation approach	Independent, held-out test data; results reproducible by a third party	Validation performed only on the team's own curated examples
Claims about "general intelligence"	Precise about the specific task the system solves	Frames a narrow tool as a step toward general problem-solving without evidence
Compute honesty	Clear about training and inference cost, and how it scales with use	Vague or evasive about the compute or data cost of running the system

Investor perspective. We are, if anything, more sceptical of AI-labelled deep-tech applications than of the other domains in this document, precisely because the label is currently the easiest one to attach to a pitch without technical substance behind it. Our practical response is to ask for the same thing every time: the specific benchmark, the specific failure mode, and who besides the team itself has reproduced the result.

DOMAIN TWO

Life sciences and biotechnology

Computational biology has produced one of the clearest, best-documented deep-tech advances of the past decade: the near-complete solution, for a large share of known proteins, of the problem of predicting a protein's three-dimensional structure from its amino-acid sequence. What that advance does and does not solve for drug discovery is a genuinely instructive case study in separating fact from promise.

State of the art

AlphaFold and its successors have predicted structures for the great majority of proteins catalogued across more than a million species, a scale of coverage that was not achievable by experimental structural biology alone within any comparable timeframe. AlphaFold 3, released in 2024, extended this to predicting how proteins interact with other molecules, including candidate drug ligands, DNA, and RNA — directly relevant to drug discovery rather than only to basic structural biology.

Technical bottlenecks that are publicly documented

- **Static structure versus dynamic behaviour.** Published technical reviews are explicit that AlphaFold's structural predictions represent a single, static conformation, without solvent, ion, or dynamic conformational information — while real proteins move and change shape, and that behaviour often matters for how a drug actually binds.
- **Wet-lab validation remains mandatory.** Reviews in the drug-discovery literature are consistent on this point: an AI-designed candidate reduces the number of variants that need to be tested experimentally, but does not eliminate the requirement for laboratory validation, animal studies, and — for a therapeutic — clinical trials, each of which remains a slow, expensive, and high-attrition stage regardless of how the candidate was generated.
- **Manufacturing and regulatory timelines are unchanged by computational speed.** A faster computational design cycle does not shorten the biological and regulatory timelines — cell-line development, toxicology, clinical trial phases — that determine when, or whether, a therapeutic actually reaches patients.

A plausible maturity horizon

Computational protein design and structure prediction are mature, widely deployed research tools today. Therapeutics that originate from an AI-assisted design process still follow the multi-year, multi-phase clinical and regulatory pathway that governs any drug candidate; nothing in the published literature suggests that pathway is being meaningfully shortened by computational front-loading, only that more candidates can be screened before that pathway begins.

> **Investor perspective.** This is, in our view, the clearest illustration in this document of why the distinction between fact and promise matters. It is an established fact that AI has meaningfully accelerated the front end of biological discovery. It is not an established fact — and we would treat any pitch that implies otherwise as a warning sign — that this acceleration has yet been shown to shorten the time or cost of bringing a therapeutic to market, because the clinical and regulatory stages downstream have not changed pace.

What distinguishes a credible team from one that oversells

- **Credible teams talk about attrition, not just hits.** They can describe how many computationally generated candidates were tested experimentally, and what fraction survived each stage, rather than only the ones that worked.
- **Credible teams are specific about which stage they are at.** Discovery, preclinical, or clinical are different categories with different risk profiles and different capital needs; a team that blurs the distinction is either inexperienced or being imprecise on purpose.
- **Credible teams do not claim regulatory approval is a formality.** Regulatory review exists because biological safety and efficacy questions are genuinely hard to answer in advance; a team that treats it as a rubber stamp misunderstands its own risk.

Why biotechnology fits a long-horizon programme

The gap between a computational discovery and a validated biological product routinely spans years, independent of how good the initial computational result was. That gap is exactly the kind of timeline mismatch described in Part Two: a discovery can be real and important, and still take longer to become a product than a conventional venture fund's typical holding period comfortably allows.

DOMAIN THREE

Energy and storage

Energy storage and generation is a domain where the gap between "demonstrated in a laboratory" and "manufacturable at scale" has repeatedly proven wider than public announcements suggested. We treat that history as a documented pattern to be taken seriously, not as cynicism.

State of the art — solid-state batteries

Solid-state batteries replace a lithium-ion cell's liquid electrolyte with a solid lithium-ion conductor, promising higher energy density and improved safety. As of the most recent published technical reviews, polymer-electrolyte concepts have reached niche commercial markets, but industrial-scale manufacturing for higher-performance chemistries — sulfide and oxide solid electrolytes in particular — still faces documented hurdles in electrolyte selection, electrode interface stability, and manufacturing process design. Industry roadmaps published by battery manufacturers describe staged commercialisation extending from consumer electronics toward automotive applications later in the decade, though the sector's own public history includes repeated schedule slippage on comparable announcements.

State of the art — nuclear fusion

Fusion is the domain's longest-horizon case, and its own technical history should be read as a caution against precise dates. In December 2022, the U.S. National Ignition Facility achieved the first laboratory demonstration of fusion ignition producing more energy from the fusion reaction than the energy delivered to the fuel target — a genuine scientific milestone. Independent analysis is equally clear, however, that this was a scientific net-energy gain measured at the target, not an engineering or economic one: the facility's lasers required roughly a hundred times more input energy to operate than the fusion reaction released, and no comparable facility has yet demonstrated the additional step — engineering gain across the whole plant — that a power-generating fusion reactor requires. Private fusion investment grew from roughly 1.9 billion US dollars in 2021 to roughly 10.5 billion US dollars by 2025 according to industry-association tracking, and several private companies have stated internal targets for pilot plants before 2030, but independent policy analysis, including work published by the University of Pennsylvania's Kleinman Center, treats grid-connected commercial fusion on that timeline as uncertain and dependent on public–private cooperation that has not yet been secured at the necessary scale.

> A specific historical caution on this domain: solid-state battery commercialisation has been publicly announced as imminent by multiple manufacturers across several successive years without matching that timeline in mass-market delivery, and fusion has a multi-decade public history of announced timelines that slipped. We treat both facts as reasons for particular care in this domain, not as reasons to dismiss it.

Technical bottlenecks that are publicly documented

Technology	Core technical bottleneck	Publicly stated horizon
Sulfide solid-state batteries	Electrolyte-electrode interface stability; lithium sulfide synthesis cost	Manufacturers' roadmaps cite commercialisation around 2030, with partial breakthroughs possible 2027–2028
Oxide solid-state batteries	Processing temperature and manufacturability at scale	Reported as earlier-stage than sulfide chemistries in current technical reviews
Magnetic-confinement fusion	Sustained plasma confinement; engineering gain across the full plant, not just the reaction	Private pilot-plant targets cluster around 2030; independent analysts treat grid delivery on that date as uncertain
Inertial-confinement fusion	Laser efficiency; repetition rate needed for continuous power generation	No public engineering-gain milestone has yet been announced

What distinguishes a credible team from one that oversells

- **Credible teams distinguish scientific gain from engineering gain** — the difference between "the reaction released more energy than the fuel received" and "the whole plant produces more energy than it consumes," which are very different claims.
- **Credible teams cite their own test data, not only vendor or press announcements**, and are specific about cycle counts, energy density measured under realistic conditions, and failure modes observed.
- **Credible teams acknowledge the sector's track record on timelines** rather than presenting their own roadmap as immune to the delays that have affected comparable programmes.

> **Investor perspective.** Energy and storage is, in our assessment, the domain in this document where the distance between a genuinely important laboratory result and a commercially meaningful outcome has most often been underestimated by the sector itself. We do not think that makes the domain uninvestable — the scientific progress described above is real — but it does mean we weigh a team's honesty about its own timeline risk more heavily here than in most other domains.

DOMAIN FOUR

Advanced materials

Materials science is undergoing the same kind of computational acceleration as structural biology, and it offers an equally instructive lesson about the gap between discovering a candidate material and being able to manufacture it.

State of the art

In 2023, Google DeepMind published GNoME, a deep-learning system that identified roughly 2.2 million candidate inorganic crystal structures, of which a large subset were assessed as stable by computational criteria — described by the researchers as equivalent to centuries of prior discovery effort by traditional means. Independent laboratories, including Lawrence Berkeley National Laboratory's autonomous "A-Lab," have gone further and used robotics combined with machine-learning guidance to synthesise some of these computationally predicted materials without a human directly running each experiment, reported in a peer-reviewed 2023 Nature paper.

Technical bottlenecks that are publicly documented

- **Zero-temperature predictions versus real synthesis conditions.** Subsequent published critical evaluation of the GNoME dataset notes that its stability predictions rely on density-functional-theory calculations performed at zero kelvin, which can favour perfectly ordered atomic structures that do not necessarily persist once a material is actually synthesised at real, elevated processing temperatures — meaning a computationally "stable" material is not automatically a synthesisable one.
- **Synthesis is still the bottleneck, not discovery.** Of millions of computationally predicted candidates, only a comparatively small number have been independently synthesised and verified in a laboratory to date; the rate-limiting step in materials innovation has shifted from finding candidates to making and testing them.
- **Scale-up from bench synthesis to manufacturable quantity** introduces its own separate engineering problems — purity control, cost of precursor materials, process repeatability — that a computational prediction says nothing about.

A plausible maturity horizon

Computational materials discovery is a mature and rapidly improving research tool today, materially expanding the space of candidates a team can consider. The translation of any individual candidate into a manufacturable, cost-competitive product remains, on the public record, an experimental and engineering process measured in years per material family, largely independent of how quickly it was identified computationally.

Investor perspective. We think this domain is a close analogue to life sciences: a genuinely powerful discovery engine sitting upstream of a slow, expensive, largely unchanged validation and scale-up process. A materials team that leads with the size of its computational search space, rather than with what fraction of its candidates has actually been synthesised and characterised, is telling us more about its marketing than about its science.

What distinguishes a credible team from one that oversells

Signal	Credible team	Warning sign
How discovery scale is framed	States how many candidates were synthesised and characterised, not only predicted	Leads with the total number of computationally predicted candidates as if it were the achievement
Synthesis conditions	Specific about temperature, pressure, and process used to produce real samples	Vague about whether a "discovered" material has ever been physically made
Scale-up plan	Names the specific manufacturing bottleneck it expects to face next	Treats scale-up as a formality once discovery is complete

Why this domain fits a long-horizon programme

The multi-year gap between a computationally predicted material and a qualified, manufacturable product is structurally similar to the biotechnology case: a real acceleration at the discovery stage, and an largely unchanged, capital-and-time-intensive validation stage downstream. A financing structure that only rewards fast validation will systematically underfund exactly the materials programmes most likely to matter.

DOMAIN FIVE

Computing and infrastructure

Every domain above depends, to some degree, on the availability of computation. This domain looks directly at the physical and architectural constraints on computing itself — both the classical infrastructure straining under AI's growth, and quantum computing as a genuinely different, much longer-horizon computational paradigm.

State of the art — classical compute constraints

The scaling pattern behind recent AI progress — roughly five-month doubling of training compute, according to Stanford's 2025 AI Index — has a direct physical counterpart in energy use, which the same report documents as growing on a roughly annual doubling cycle. This is now widely treated in the technical literature as a genuine infrastructure constraint: power availability and delivery, not just chip supply, increasingly shapes where and how fast frontier computing capacity can be built.

State of the art — quantum computing

Quantum computers process information in ways that could, for a specific and still relatively narrow class of problems, outperform classical computers by exploiting quantum superposition and entanglement. The central open technical problem is error correction: individual physical qubits are highly error-prone, so useful computation requires encoding one reliable "logical" qubit across many physical qubits — current industry estimates cited by quantum-hardware researchers put this at roughly 100 to 1,000 physical qubits per logical qubit, depending on the error-correction code and hardware platform used. Published research volume in this specific area grew sharply through 2025, with peer-reviewed papers on quantum error-correction codes rising from 36 in all of 2024 to 120 in the first ten months of 2025 alone, and major hardware developers including IBM and Google have published explicit multi-year roadmaps toward large-scale fault-tolerant machines, each built around specific, named error-correction code families.

Technical bottlenecks that are publicly documented

- **Physical-to-logical qubit overhead.** The 100-to-1,000-to-one ratio between physical and logical qubits means that a fault-tolerant machine capable of genuinely useful computation requires physical qubit counts far beyond what any current public demonstration has achieved.
- **Real-time decoding.** Error-correction codes are only useful if the classical computation needed to detect and correct errors can run fast enough to keep up with the quantum hardware — a distinct and actively researched engineering problem in its own right, separate from qubit count.
- **Energy delivery for classical AI infrastructure.** Annual doubling of power consumption for frontier AI training and inference is now a documented planning constraint for grid

operators and data-centre developers, not a hypothetical future concern.

A plausible maturity horizon

Both hardware developers publishing public roadmaps in this space — IBM and Google among them — describe multi-year, staged paths toward large-scale fault-tolerant quantum computing, each organised around specific technical milestones (logical qubit counts, error rates, decoding speed) rather than calendar dates alone. The sharp increase in peer-reviewed error-correction research through 2025 is a genuine and measurable sign of momentum. It remains, on every public roadmap we have reviewed, a multi-year engineering programme rather than a near-term product category, and classical-versus-quantum performance comparisons remain contested for all but a narrow set of specialised problems.

What distinguishes a credible team from one that oversells

- **Credible teams state their logical-qubit count, not only their physical-qubit count** — the two numbers can differ by two to three orders of magnitude, and conflating them is one of the most common ways this field is oversold publicly.
- **Credible teams name a specific error-correction code and cite a measured logical error rate**, rather than describing progress only in qualitative terms.
- **Credible infrastructure teams are specific about power and cooling constraints**, since these are now binding constraints on classical AI infrastructure just as much as chip availability.

> **Investor perspective.** Quantum computing is, in our view, the domain in this document furthest from a near-term commercial inflection point, and we say so directly rather than softening it. That does not make it uninvestable on a genuinely long horizon — the underlying research momentum is real and measurable — but a ten-year holding horizon, which we discuss directly in Part Six, is arguably closer to necessary than optional for this specific domain, more than for any other covered here.

PART FOUR — THE EVALUATION METHOD

How we judge a scientific project before it has a market

Evaluating a deep-tech opportunity before commercial traction exists requires different questions than evaluating a business with a revenue line. The questions below are the ones we actually put to a team, organised by what they are designed to reveal.

Questions about the science itself

- What, specifically, has been demonstrated — measured, published, or independently reproduced — as opposed to modelled, simulated, or projected?
- What is the single most likely technical reason this could fail, in the team's own words, and how would they know if it was failing?
- What would a domain expert with no stake in this team's success say is the weakest part of the claim?
- Which specific published state-of-the-art result is this meant to improve on, and by how much, measured how?

Questions about the team

- Who on the team has done the specific technical work being proposed before, in what setting, with what outcome — including outcomes that did not work?
- Can the team describe their own prior failed experiments as clearly as their successes?
- Is the team's technical leadership able to explain the work directly, or only through a commercially framed intermediary?

Questions about the path ahead

- What is the next specific, falsifiable milestone, and what would count as clearly not meeting it?
- What does the team need to learn, not just build, in the next phase of work?
- What existing infrastructure, equipment, or specialised capability does this depend on, and how exposed is the plan to losing access to it?

Warning signals we weigh heavily

Signal	Why it concerns us
No discussion of failure modes or limitations, unprompted	Every real technical programme we have reviewed has known weaknesses; their absence from a pitch usually means they were omitted, not that they don't exist
Validation performed only on data the team itself curated	Self-selected validation is the single easiest way to make a weak result look strong
Comparisons to "AlphaFold-like" or "GPT-like" breakthroughs without a specific, checkable benchmark	Borrowing the credibility of a well-known result without submitting to the same standard of evidence
Inability to name the specific published state of the art being improved on	Suggests the team has not engaged seriously with existing published work in its own field
Timeline that does not account for the field's own documented history of delay	As seen in energy storage and fusion above, ignoring a sector's own track record is a specific, checkable form of overselling

Stages of validation we look for

We do not expect every stage below to be complete at the point we engage — for an early-stage team, most will not be. What we look for is a team that can locate itself honestly on this progression, and that understands what each subsequent stage will actually require.

1. **Internal proof of concept** — a result reproduced reliably inside the team's own lab.
2. **Independent reproduction** — the same result, or a closely related one, reproduced by a party with no stake in the outcome.
3. **Scaled or repeated demonstration** — the result holding up outside the narrow conditions of the first demonstration.
4. **Path to manufacture or deployment defined** — a credible, specific description of what stands between the validated result and a usable product, not just an assumption that it will follow.

Investor perspective. None of these questions is designed to catch a team lying. Most overselling we encounter in this category is not dishonest — it is founders who have spent years believing in their own work, understandably, and have lost the outside perspective needed to state its limits plainly. Our method is built to supply that outside perspective early, before it becomes expensive to discover later.

PART FIVE — THE SUPPORT MODEL

What patient capital brings beyond money, and how we accept scientific failure

Capital alone rarely determines whether a hard-science programme succeeds. What surrounds the capital — how milestones are set, how governance behaves under uncertainty, and how a negative result is handled — shapes whether a team can do its best work.

Beyond the cheque

Consistent with how Future Power Capital operates across every activity described in our Firm Overview, we act as principal and bring the same practical operating judgment to a Deep Tech Programme commitment that we bring elsewhere — organisational structure, hiring sequencing, and the discipline of preparing for outside technical or commercial scrutiny before it arrives. For a research-led team, that often means help thinking through questions that are adjacent to the science but not the science itself: how to sequence hires between research and engineering, how to structure a relationship with a specialised supplier or facility, and how to prepare technical results for an audience that will not read them the way the team's own field does.

How we think about milestones

We set milestones around what a field itself recognises as meaningful progress — the validation stages described in Part Four — rather than around a fundraising or fund-lifecycle calendar that has nothing to do with the science. A milestone in this framework is a question to be answered, not a number to be hit: "does the mechanism work under realistic conditions," not "did revenue grow by a fixed percentage this quarter."

> This describes our intended approach to a long-horizon programme under active development, as stated on our Deep Tech Programme page. It is not a description of a specific current arrangement with any named team, and none should be inferred.

Accepting scientific failure honestly

A meaningful share of genuinely ambitious scientific work will not succeed — that is a structural feature of research at the frontier, not a sign of a badly run programme. We think there is an important difference between two kinds of negative result, and we intend to treat them differently:

- **A rigorous negative result** — the team asked a well-posed question, ran a well-instrumented experiment, and the answer was no. This is often as valuable, scientifically, as a positive result, because it closes off a path other people might otherwise waste years exploring.
- **A negative outcome caused by poor execution, dishonesty about progress, or an unwillingness to test the hardest version of the hypothesis** — a different category, and one that speaks to the team rather than only to the science.

The evaluation method set out in Part Four is designed to make this distinction visible while a programme is running, not only after it has concluded — by insisting on falsifiable milestones and independent validation from the outset, a rigorous failure and a poorly executed one leave a different, checkable trail.

What we mean, and do not mean, by "accompaniment"

We do not present ourselves as scientific collaborators, co-inventors, or a substitute for peer review. Our role is to fund responsibly, ask honest and informed questions, help with the organisational and commercial dimensions of running a technical company, and give a team the time a genuinely hard problem requires — not to direct the research itself.

> **Investor perspective.** We think a financing relationship that punishes every negative result equally, regardless of whether the underlying work was rigorous, quietly teaches teams to hide problems rather than surface them early — which is worse for everyone, including the capital involved. Distinguishing a well-run failure from a poorly run one is, in our view, one of the more important and least discussed parts of funding hard science responsibly.

PART SIX — THE TEN-YEAR HORIZON

What ten years actually means, in milestones rather than slogans

"Long horizon" is easy to state and easy to leave vague. We think it is more honest to describe, concretely, what a decade-scale relationship with a research-led team can look like in stages — while being clear that not every programme will move through every stage, and some will end earlier.

An illustrative decade, not a guarantee

The stages below are illustrative of how a long-horizon relationship can unfold, drawn from the general shape of technical development described across the domains in Part Three. They are not a timeline for any specific programme, current or future, and durations will vary — often considerably — by domain, as the maturity horizons in Part Three make clear.

Phase	What it typically involves	What "progress" looks like
Years 0–2	Core technical validation; the team answers the hardest open question in its own approach	A rigorous internal result, ideally independently reproduced
Years 2–4	Moving from a single demonstrated result to a repeatable process or a scaled prototype	The result holds outside the narrow conditions of the first demonstration
Years 4–7	Engineering and manufacturing realities — the gap described repeatedly in Part Three between a validated result and a producible one	A credible, tested path to production, not just a plan for one
Years 7–10	Commercial or scientific deployment at a scale that tests the result against the real world outside a controlled setting	The technology performing under conditions the team does not fully control

> Every domain in Part Three sits somewhere different on this scale today. Computational drug and materials design already operate inside the years-2-to-4 range on a mature basis; large-scale fault-tolerant quantum computing and grid-connected commercial fusion, on public roadmaps, are closer to the years-7-to-10 range or beyond it. A ten-year horizon is not one clock that starts at the same point for every domain.

Why value builds across this duration, not despite it

The financing gap described in Part Two exists specifically at the stage where technical risk is still real but the capital required to keep testing it has grown — the "missing middle" documented in deep-tech financing research. A structure that can hold a position through that stage, rather than needing to exit before it or refusing to enter until after it, is positioned differently from capital that cannot. We make no claim about what that positioning is worth in return terms — we make none anywhere in this document — but we think the structural logic is sound: capital that can survive the stage where most financing structures cannot is, by definition, present for a different part of a technology's development than capital that cannot.

What we watch for across the decade

- Whether each phase transition in the table above is supported by the kind of independent validation described in Part Four, not just by the team's own account of progress.
- Whether the team's own stated timeline honestly reflects the domain's documented history — as Part Three shows, some fields have a public track record of schedule slippage that a credible team should acknowledge, not ignore.
- Whether the reasons we backed the team originally still hold, particularly after a negative or ambiguous result, consistent with how we describe accepting scientific failure in Part Five.

> **Investor perspective.** We think the honest version of "ten-year horizon" is not a promise that ten years guarantees an outcome — it does not, and nothing in this document should be read as suggesting otherwise. It is a statement that some of the most important scientific and engineering problems, on the public record summarised in Part Three, simply do not resolve on a three-to-five-year fund cycle, and that capital unwilling to operate on a longer clock will structurally miss a share of the work most worth funding.

CONCLUSION

A closing note, and an invitation to research teams

The domains covered in this document sit at very different points on the path from laboratory result to real-world outcome, and we have tried throughout to be plain about which is which. Applied artificial intelligence and computational methods in biology and materials science are already powerful, deployed research tools, sitting upstream of validation and manufacturing stages that remain slow and expensive on their own terms. Energy storage and fusion carry a documented history of optimistic public timelines that this document has not tried to soften. Quantum computing is, on its own public roadmaps, the furthest from a near-term commercial inflection point of everything covered here.

We think a long-horizon, independently governed capital structure has a genuine, structural role to play across this range — not because it improves the odds of any individual project, which we make no claim about, but because it removes a specific, well-documented mismatch between how research timelines actually unfold and how most financing structures are built to behave.

What this means in practice, today

The Deep Tech Programme, as our services page states plainly, is a long-horizon initiative under active development. This document is the fullest account, at the time of writing, of the reasoning behind it — the thesis, the financing problem it responds to, the domains it intends to cover, the method it intends to apply, and the kind of decade-scale relationship it is built to support.

An invitation to research-led teams

If your work sits in the gap described in Part Two — too early or too capital-intensive for conventional venture capital, and too commercially oriented for a typical grant cycle — and you can describe your own science honestly, including what remains unproven, we would like to hear from you through the process described on our website. We are not able to promise funding, a specific timeline for a response, or any outcome of a conversation; what we can promise is that we will engage with the substance of the science on its own terms, using the method set out in Part Four of this document.

> *The problem worth funding is rarely the one that is easiest to explain in a single sentence. We are built to sit with the harder ones for as long as they actually take.*

Get in touch

EMAIL

contact@futurepower-capital.com

TELEPHONE

+44 20 3355 8354

REGISTERED OFFICE

The Tea Building, 56 Shoreditch High Street, London, E1 6JJ, United Kingdom

FOUNDER & DIRECTOR

Silas Trent

Future Power Capital Ltd — company number 16778446 (England & Wales). Registered office: The Tea Building, 56 Shoreditch High Street, London, E1 6JJ, United Kingdom. Incorporated 10 October 2025. Founder & Director: Silas Trent.

Issued and fully paid share capital: GBP 10,000,000 (10,000,000 ordinary shares of GBP 1 each).

View our official record on Companies House