ARTIFICIAL ENERGY

THE CIVILIZATIONAL PIVOT FROM DESTROYING MATTER TO HARVESTING GRADIENTS

SAME PHYSICS. DIFFERENT MATTER. REAL OUTCOME.

JED ANDERSON

Writer and builder. Working at the intersection of
physics, information theory, and the protection of nature.

2026 · Second Edition — Revised

“The living organism feeds upon negative entropy.”

—ERWIN SCHRÖDINGER, What is Life? (1944)

“It from bit.”

—JOHN ARCHIBALD WHEELER (1990)

EXECUTIVE SUMMARY

THE ARGUMENT IN TWO PAGES

The essay is not about traditional energy production. It is not about coal, oil, gas, nuclear fission, or familiar renewables like solar, wind, and hydro as competing entries on the same fuel menu. For two hundred years, the predominant way humans have produced energy has been a single operation: disassembling matter. We break chemical bonds in combustion, split heavy nuclei in fission, and now attempt to fuse light ones. In every case, the fuel and the apparatus are the same matter; the substrate is sacrificed and waste accumulates. Recent shifts toward renewables change the fuels and devices but not the category of energy they belong to. This essay steps outside that frame entirely. It is about a different regime of energy: not consuming matter, not pointing passive apparatus at a gradient, but engineering matter with enough information to harvest work from the free-energy gradients the universe already provides.

For billions of years, a planetary-scale information-structured energy architecture has been running without being recognized as a coherent system. Its working demonstration is the biosphere. The biosphere does not consume its own apparatus; it uses informationally structured matter—chlorophyll, enzymes, evolved molecular machines—to harvest gradients that already exist. We have been looking at that architecture the whole time without seeing what we were looking at.

Once that architecture is seen, it can be built in other substrates. This essay names the regime and the technology that will engineer it at industrial scale: **Artificial Energy (AE)**. The parallel to Artificial Intelligence is structural. AI is intelligence engineered in non-biological substrates (silicon, software) doing what biological brains do by the same physics in different matter. AE is gradient harvesting engineered in non-biological substrates (designed materials, engineered catalysts, artificial reaction centers) doing what biological photosynthesis and metabolism do by the same physics in different matter.

The substrate is artificial. The physics is real. The energy is real. What this names is not a better way to capture sunlight. It is access to a deeper and far larger reservoir of usable energy: the free-energy gradients that pervade the universe, harvested by matter structured with enough information to steer them. Biology has drawn on this reservoir since before there were continents. Its engineered forms—artificial photosynthesis, engineered enzymatic catalysis, AI-designed harvesting and steering systems—barely exist yet. They are the next civilizational technology layer, parallel to AI in significance, and the opportunity to build them is open now.

The reservoir is not new. It is older than life. What is new is the recognition that it can be reached deliberately, by design, at industrial scale—and that reaching it is a single coherent enterprise deserving a single name. That recognition is the whole of this essay.

I THE THING HAS BEEN MISNAMED

The single most important step in seeing the world clearly is naming things correctly. Energy has been misnamed for two hundred years.

When humans say 'energy,' they almost always mean one specific operation: disassembling matter to release the energy bound inside it. Burning a log breaks the C-H and C-C bonds in cellulose and releases roughly 5 electron-volts per bond. Splitting a uranium nucleus releases 200 million electron-volts. Fusing deuterium and tritium releases 17.6 million electron-volts. The predominant way industrial civilization has produced energy for two hundred years (coal, oil, gas, fission, and the long pursuit of fusion) follows the same recipe. Find matter with high binding energy. Disassemble it. Capture the energy that comes out.

A note on naming. Combustion is not 'mass destruction' in any strict sense. The mass converted to energy in breaking a C-H bond is on the order of four parts in ten billion. The nuclei are spectators; only electrons rearrange. Fission and fusion convert measurably larger fractions of mass (about 0.1% and 0.4% respectively), but even these are small numbers. What unifies combustion, fission, and fusion is something more important than mass-energy conversion ratios: in all three regimes, the fuel and the apparatus are the same matter. The fuel is consumed in the act of producing energy. The substrate of the reaction is sacrificed. This essay uses the phrase mass-disassembly to name the regime when precision matters, and mass-destruction when colloquial force matters. Both denote the same physical distinction: the fuel itself is the thing being taken apart.

This is one regime. There is another.

The other regime does not consume the apparatus that performs the work. It uses informationally-structured matter (molecules, lattices, devices designed with specific architectures) to couple to free-energy gradients that already exist in the universe, and to extract useful work as those gradients relax. Sunlight against the cold of space. A warm room next to a cold one. Hydrogen and oxygen that have not yet reacted. A salt front meeting fresh water. In each case, the structured apparatus does not run down. It directs the work that the gradient does on its own. The fuel and the apparatus are separated: the apparatus persists; the gradient supplies the energy.

Both regimes obey the conservation of energy. Both produce real, usable work. They are not in conflict. They are simply two different physics regimes available to any civilization with sufficient capability. The choice between them is not a question of physics. It is a question of strategy.

A second distinction must be drawn immediately, because without it the argument is easy to mistake for something familiar. Within the gradient-harvesting regime there is a wide range of sophistication, and the technologies the public already associates with non-combustion energy sit at the very bottom of it. A solar panel, a wind turbine, a hydroelectric dam, a geothermal plant: each of these harvests a gradient rather than consuming a fuel, and each is therefore on the right side of the first divide. But each does so in the crudest possible way. It couples a fixed, passive structure to a single bulk gradient that happens to be concentrated enough to exploit, the sun on a flat absorber, wind on a blade, water falling under gravity, heat conducted from hot rock, and converts what it can into electricity. There is almost no information in the apparatus. It selects nothing, routes nothing, transforms nothing at the molecular scale. It is a sieve held up to a stream.

This essay is about the move beyond that crude form, not a defense of it. The deeper opportunity is the high-information end of the regime: matter engineered with enough structure to harvest gradients selectively, at the molecular and quantum scale. A designed catalyst chooses one reaction among millions. An engineered enzyme assembles a specific molecule at ambient temperature. An artificial reaction center routes individual quanta of energy along a chosen path. These devices reach gradients that no panel or turbine can touch, chemical disequilibria, specific molecular bonds, diffuse and low-grade differences, and they produce not merely electricity but fuels, fertilizers, and designed materials, because the apparatus performs chemistry rather than bulk conversion. The distinction is not which gradient is tapped. It is how much design intelligence is built into the matter that taps it, and therefore how much of the available world it can reach. The first harvesters were the beginning of the regime. They are not its destination, and they are not the subject of this essay.

TWO DIVIDES, NOT ONE

The confusion between passive gradient energy and Artificial Energy is strategic, not merely semantic. The first divide between substrate-consuming and gradient-harvesting energy is the basis for most clean energy advocacy. The second divide, less recognized but equally important, separates passive gradient coupling from information-rich gradient steering. A solar farm and a natural-gas plant look like opposites: one burns nothing, the other burns continuously. But they share a conceptual architecture: passive apparatus pointed at a resource, producing undifferentiated output. Artificial Energy belongs to a different category entirely. Its apparatus selects, routes, and transforms. It performs chemistry. It reads the molecular environment and responds to it. Naming solar and wind 'renewable' was accurate and useful . . . it freed them from the carbon frame. Naming AE 'Artificial Energy' is the next move . . . freeing the information-rich frontier from the passive-device frame that renewable energy still inhabits.

Tier	Name	Examples	Defining Characteristic
1	**Mass-Destruction Energy**	Coal, oil, gas, nuclear fission, fusion	Substrate consumed; apparatus is the fuel; waste is intrinsic to the reaction
2	**Passive Gradient Energy**	Solar PV, wind turbines, hydroelectric dams, geothermal	Gradient harvested; apparatus is fixed and passive; near-zero information; single-mode output; bounded by device physics
3	**Artificial Energy (AE)**	Artificial photosynthesis, AI-designed catalysts, engineered enzymes, thermophotovoltaics, biomimetic membranes	Gradient harvested by information-rich engineered matter; selective, molecular-scale, multi-output; the subject of this essay

Industrial civilization, beginning around 1800, leaned overwhelmingly toward mass-disassembly. The biosphere, beginning around 3.5 billion years ago, ran entirely on gradient harvesting. The biosphere's track record is staggering at the level of system architecture, though the comparison to industrial energy requires care. Net primary productivity, the rate at which biological gradient harvesting converts sunlight into chemical bonds, runs on the order of 130 terawatts globally. This is not directly comparable to industrial civilization's roughly 20 terawatts of delivered, high-exergy power, because most biospheric throughput is low-grade chemical storage that cycles back to CO_2 within months. What the comparison does establish is that gradient harvesting is feasible at planetary scale. A system of distributed, informationally-structured matter has been operating for billions of years, at ambient temperature, with no combustion, no fission, no fusion. It is not waste-free: biology produces carbon dioxide, methane, and reactive nitrogen, and redistributes matter constantly. What distinguishes it is that a mature biosphere recycles these material flows rather than accumulating them, while its energy balance closes by radiating low-grade heat to space. The biosphere captures only about a tenth of one percent of the roughly 90,000 to 115,000 terawatts of solar exergy that reaches Earth's surface. Even at that low capture rate, the system runs.

The argument of this essay is simple, and it requires drawing two divides, not one. The first divide separates mass-destruction energy — combustion, fission, fusion — from gradient-harvesting energy. The second divide, less widely recognized but equally consequential, separates passive gradient harvesters from information-rich gradient harvesters. Solar panels, wind turbines, and hydroelectric dams belong to the middle tier: they harvest gradients rather than consuming a fuel, which is progress, but they do so with near-zero information in the apparatus. They select nothing. They route nothing. They transform nothing at the molecular scale. Most of what is now called 'clean energy' occupies this middle tier — one step beyond mass-destruction, but still operating at the low-information floor of what gradient harvesting can do. Artificial Energy names the third tier: matter engineered with enough structure to harvest gradients selectively, at the molecular and quantum scale, producing not merely electricity but fuels, fertilizers, and designed materials — the way biology has done for 3.5 billion years. Industrial civilization has been doing Tier 1 for two hundred years and is approaching its limits: climatic, geopolitical, thermodynamic. The biosphere has been doing Tier 3 since before there were continents. The next civilizational technology layer is to engineer Tier 3 deliberately. The name proposed for it here is Artificial Energy, and the case for that name rests on first principles.

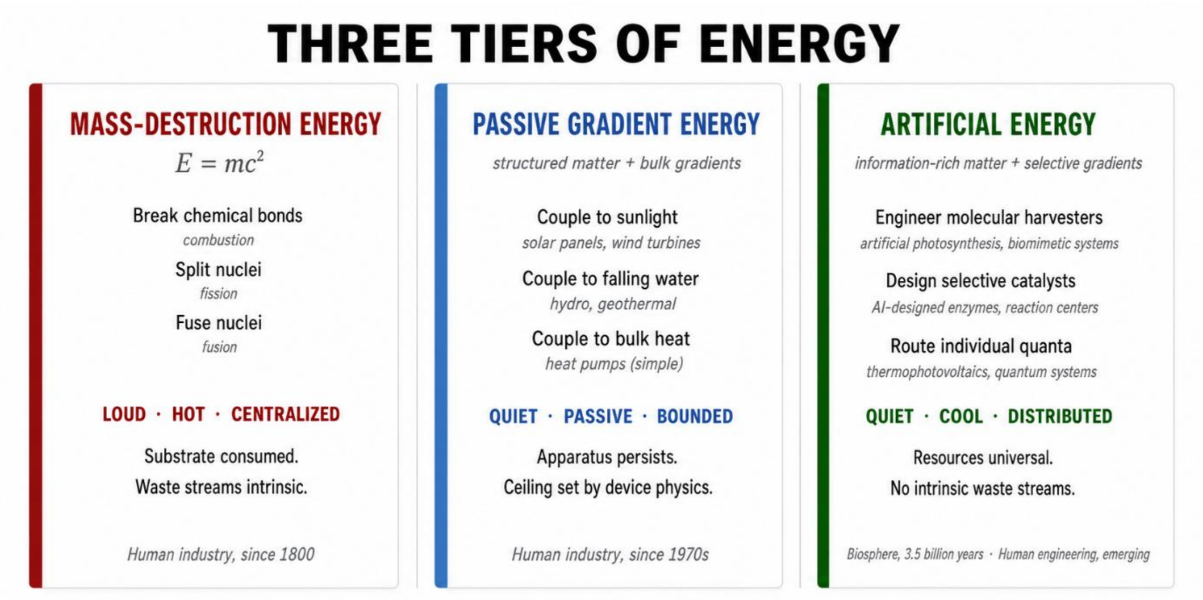

Figure 1. The three energy tiers. Tier 1 (Mass-Destruction): combustion, fission, fusion — substrate consumed, apparatus sacrificed, waste intrinsic. Tier 2 (Passive Gradient Energy): solar panels, wind turbines, hydro — gradient harvested, but apparatus is fixed, passive, and near-zero-information. Tier 3 (Artificial Energy): information-rich engineered matter harvesting gradients selectively at molecular scale — the subject of this essay. Both gradient tiers obey conservation laws; only Tier 3 can scale without poisoning the host.

II ARTIFICIAL INTELLIGENCE IS NOT ARTIFICIAL

To understand why Artificial Energy is the right name requires first understanding why Artificial Intelligence is the right name, and why, on closer inspection, it is also the wrong one.

The phrase 'artificial intelligence' was coined by John McCarthy at the 1956 Dartmouth Workshop, and it has done seventy years of heavy semantic lifting. In its everyday usage, 'artificial' suggests fake, imitation, second-best: plastic flowers, artificial sweeteners, the artificial smile of someone who would rather be elsewhere. By this connotation, AI is something pretending to be intelligent without truly being intelligent. This is the framing in popular discourse, and it is technically wrong.

The word 'artificial' descends from the Latin artificialis, meaning 'made by skill or art.' The root sense is not 'fake.' The root sense is 'made through human craft rather than arising spontaneously in nature.' An artificial heart is a real heart performing real cardiac function in an engineered substrate. An artificial limb performs real motor work. An artificial reef shelters real fish. The 'artificial' modifier denotes the origin of the apparatus, not the reality of the function.

By this proper reading, Artificial Intelligence is intelligence, actual and functional and measurable intelligence, engineered in a non-biological substrate. Silicon, software, designed architectures rather than evolved ones. AI does the same kind of thing biological intelligence does: structured information processing under uncertainty, by the same underlying physics, in different matter. The information it processes is real. The pattern-recognition it performs is real. The capabilities it produces are real.

In a deeper sense, then, Artificial Intelligence is not artificial at all. Intelligence is intelligence, whatever the substrate. A pocket calculator was not metaphorically performing arithmetic. ChatGPT is not metaphorically processing language. The 'artificial' word is a useful flag for substrate origin, but it should not be allowed to imply that the function is somehow lesser. Intelligence is what intelligence does.

Biological neurons and silicon transistors are two implementations of the same underlying class of phenomena.

> ***Artificial Intelligence is not artificial. It is intelligence in a different matter. The word “artificial” tells you where the apparatus came from, not whether the capability is real.***

This recognition unlocks the parallel. If AI is intelligence engineered in non-biological substrates, silicon doing what neurons do by the same physics in different matter, then there is an obvious second member of the genus. Energy capture, the other thing biological matter does in industrial volumes, is also being engineered in non-biological substrates. Silicon photovoltaics, engineered catalysts, biomimetic reaction centers, AI-designed harvesting systems. Doing what chlorophyll and enzymes do, by the same physics, in different matter.

This is Artificial Energy. The same etymology applies. The same defense holds. The energy is real. The harvesting is real. The substrate is engineered. The word 'artificial' tells you where the apparatus came from, not whether the energy that comes out is fake. It is not fake. It is identical, electron-volt for electron-volt, to the energy a leaf produces. The leaf evolved its apparatus. Engineering produced the other. That is the only difference.

One boundary must be drawn at the outset, because it defines what this essay is about and what it is not. Gradient harvesting is a wide regime, and not every member of it is the subject here. The defining axis of Artificial Energy is the amount of information built into the harvesting apparatus. A bare silicon panel and a wind rotor sit at the low-information floor of the regime: each couples to a single gradient in a single fixed way, with almost no selectivity, and each runs quickly into a hard physical ceiling. They prove the regime is real, and they are not the frontier. The frontier is the high-information end: designed catalysts, engineered enzymes, artificial reaction centers, and adaptive materials that selectively route, transform, and combine energy flows the way a chlorophyll complex or a metabolic pathway does. When this essay speaks of Artificial Energy as the next civilizational layer, it means the information-rich end of the regime, the part that barely exists yet, not the panels and turbines that have already scaled. The variable that matters is not which gradient is tapped, but how much design intelligence is embedded in the matter that taps it.

III THE ENERGY LADDER

Every energy event in the universe sits somewhere on a single ladder of magnitudes. The ladder organizes the entire argument.

The unit appropriate to single events is the electron-volt: the energy required to move one electron through a potential difference of one volt. In this unit, the entire energy hierarchy fits in one column. The numbers span eleven orders of magnitude. From the bit to the bond is a factor of at least 240, the bond-bit ratio. From the bond to the nuclear event is a factor of several million.

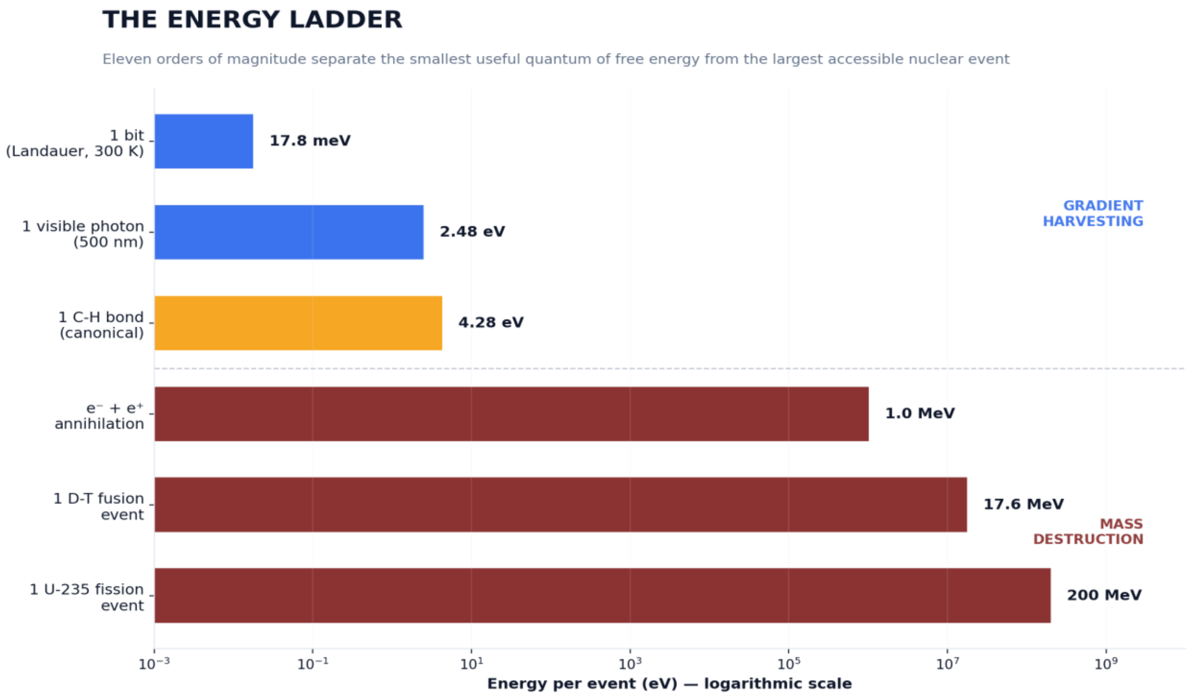

Figure 2. Energy released per single event, on a logarithmic scale. The lower three bars are gradient-harvesting territory; the upper three are mass destruction. A C-H bond, a single fission event, and a single fusion event are separated by roughly a million-fold each.

What governs these magnitudes is the fraction of rest mass that is converted to energy in each event. Chemistry rearranges electrons; the nuclei are spectators and their mass is essentially preserved; fewer than four parts in ten billion is converted. Nuclear reactions rearrange nucleons inside the nucleus itself; a measurable fraction of mass (about 0.1% for fission, 0.4% for fusion) disappears as binding energy. Matter-antimatter annihilation converts 100% of the participating rest mass. The deeper the disassembly, the larger the mass fraction released.

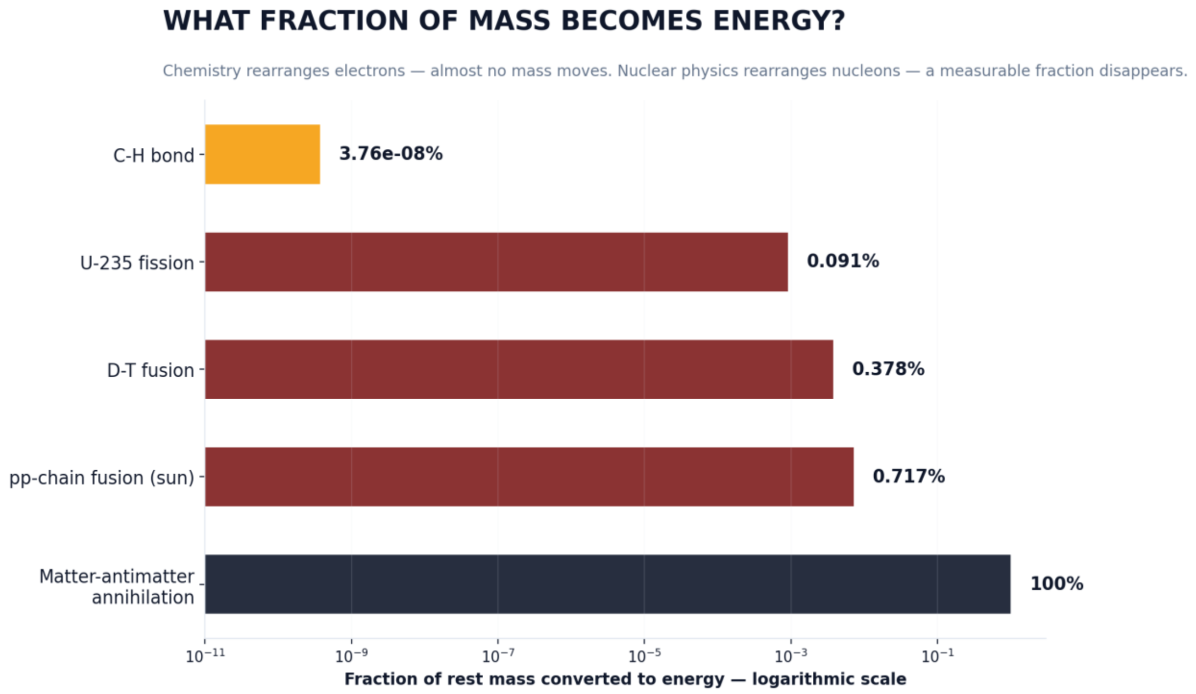

Figure 3. Fraction of rest mass converted to energy. Chemistry touches the skin of the atom. Nuclear physics touches the spine. Annihilation touches everything.

Where information lives on the ladder

At the bottom of the ladder, far below chemistry, sits the information event. Information has no intrinsic energy; a bit is a logical state, and the same bit can ride on any physical carrier. What physics fixes is the cost of handling it irreversibly. Landauer's bound states that erasing one bit must dissipate at least $kT \ln 2$ of heat, approximately 0.018 electron-volts, or 18 millielectron-volts, at room temperature. The same quantity sets the ceiling on the work an idealized information engine can extract per bit it consumes, under full thermodynamic bookkeeping. It follows from a single-line derivation: an isothermal expansion of one molecule from volume $V/2$ to V at temperature T yields work $W = \int(V/2 \text{ to } V)\ (kT/V)\ dV = kT \ln 2$. This is a bound on the cost of processing information, not a claim that a bit contains energy.

A bit is therefore at least 240 times poorer in energy content than a chemical bond, and roughly 140 times poorer than a single visible photon. This factor of 240 is the bond-bit ratio: the floor of the asymmetry between the cost of handling information and the cost of breaking matter, fixed by the second law of thermodynamics and derived in full in the canonical source (Anderson, The Bond-Bit Ratio, 2026). Information is not a richer fuel than chemistry. It is the poorest useful quantum of free energy that physics permits. What distinguishes information is not how much energy it carries. It is how cheap it is to acquire, copy, and direct, and how completely it can steer the flow of much larger energies. This ratio, and its much larger operational form, returns below in the quantification of Artificial Energy's leverage.

FIRST-PRINCIPLES VERIFICATION

Per single C-H bond: 6.86×10^{-19} J = 4.28 eV (canonical reference, 413 kJ/mol)

Per single visible photon: 3.97×10^{-19} J = 2.48 eV (500 nm)

Per single fission event: 3.20×10^{-11} J = 200 MeV (≈ 47 million × a bond)

Per single fusion event: 2.82×10^{-12} J = 17.6 MeV (≈ 4 million × a bond)

Per single bit at 300 K: 2.87×10^{-21} J = 0.018 eV (1/240 of a bond: the bond-bit ratio)

All values derive from first principles. The C-H bond enthalpy and the Landauer bound at 300 K give the canonical bond-bit ratio of approximately 240 (Anderson, 2026). Nuclear values follow from the binding-energy curve and have been validated against decades of experimental data. The Landauer bound has been experimentally verified to within 10% at the lab scale (Bérut et al., 2012).

IV WHAT THE EXPERIMENTS PROVED

Information is not metaphorically connected to energy. It is thermodynamically connected to energy. This has been measured.

For almost a century after James Clerk Maxwell proposed his demon in 1867, and after Leo Szilard formalized the engine in 1929, the relationship between information and thermodynamics was philosophical. Then it was theoretical. Beginning in 2010, it became experimental.

Toyabe et al., 2010 (Nature Physics)

Shoichi Toyabe and colleagues constructed a microscopic spiral electric potential and placed a Brownian-motion bead inside it. They measured the bead's thermal position and, based on that information, applied feedback that ratcheted the bead uphill against the potential. The bead acquired approximately 0.28 kT of free energy per measurement, with no energy input to the bead itself. The work came from the surrounding heat bath; the information selected which thermal fluctuations to harvest. This was the first direct experimental demonstration that information acquisition converts thermal noise into useful work.

Koski et al., 2014 (PNAS)

A team at Aalto University in Finland built an autonomous Maxwell's demon from single-electron transistors. The demon monitored electron positions in real time and applied feedback that extracted approximately kT ln 2 of work per bit of information acquired, within 10% of the theoretical Landauer ceiling. The experiment ran continuously, not as a single-shot demonstration. The authors described their device as the first realization of a Szilard engine extracting nearly kT ln 2 of work for one bit of information.

Ribezzi-Crivellari & Ritort, 2019 (Nature Physics)

Researchers at the University of Barcelona ran a continuous Maxwell's demon on a single DNA hairpin, repeatedly measuring whether the molecule was folded or unfolded and extracting work each time it changed state. They reported the device 'capable of extracting arbitrarily large amounts of work per cycle by repeated measurements of the state of a system,' achieving approximately 90 percent of the Landauer ceiling per bit and operating in a sustained continuous mode.

The three experiments together establish, beyond reasonable dispute, that information is a real thermodynamic resource. Per bit, the work extractable is bounded above by kT ln 2: small compared to a chemical bond, but real, measured, and indistinguishable in physical character from the work that any heat engine produces. The Second Law is preserved by the bookkeeping: every bit acquired must eventually be erased, and erasure costs at least kT ln 2 of heat dissipation. But acquisition and erasure can be separated. Memory tape, freshly blank, is itself a thermodynamic resource. Charles Bennett showed in 1982 that a blank tape is formally equivalent to a cold reservoir.

> ***Information is not fuel. It carries no energy of its own. What it does is direct the harvest of energy that already exists, at a handling cost far below the energies it steers. That is not a smaller fuel. It is a different role entirely.***

V THE BIOSPHERE'S 3.5-BILLION-YEAR DEMONSTRATION

If gradient harvesting via informationally-structured matter were difficult, slow, or limited in scale, biology would not exist.

The strongest evidence that gradient harvesting is a mature regime is that an entire planetary-scale system has been running on it, continuously, for approximately 3.5 billion years. The biosphere processes about 130 terawatts of free energy in net primary productivity, converting sunlight into chemical bonds at a rate that is not directly comparable to industrial electricity but is unambiguously planetary in scope. It does this entirely at ambient temperature, with no combustion, no fission, no fusion. It is not waste-free, but it recycles its material flows rather than accumulating them, and its energy balance closes by radiating low-grade heat to space. The biosphere captures only about a tenth of one percent of the available solar exergy reaching Earth's surface; the relevant point is not that biology operates at high efficiency, but that the harvesting regime works at planetary scale at all. No civilization had to invent it; evolution found it billions of years before humans existed.

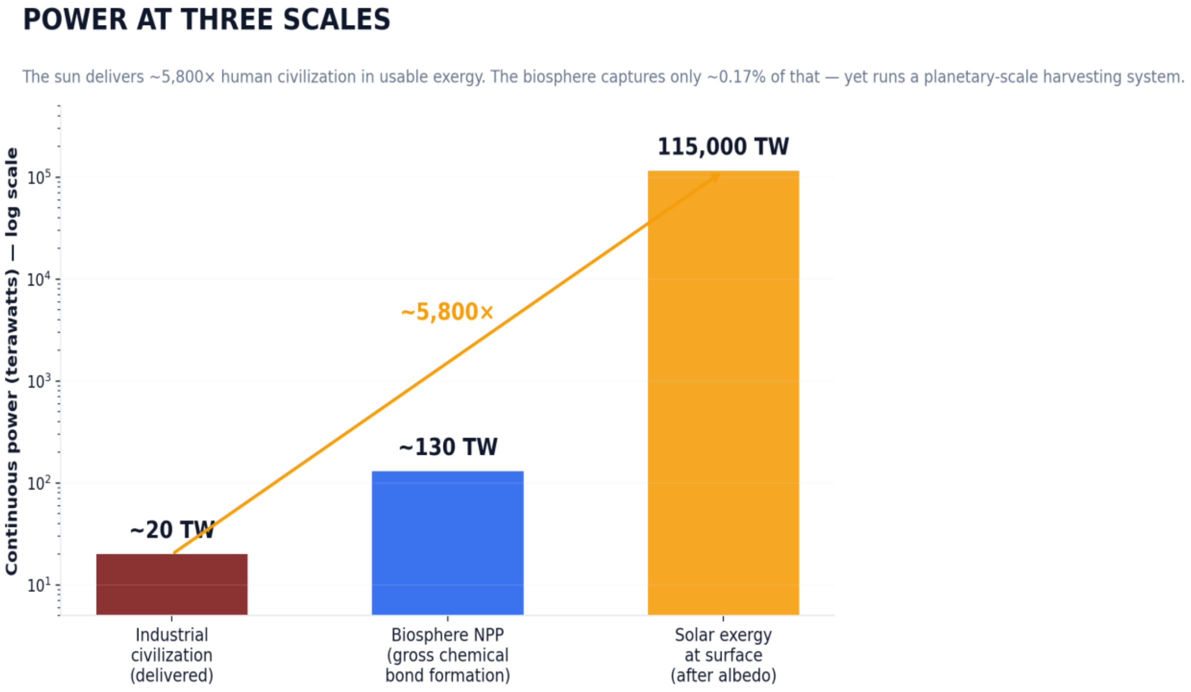
POWER AT THREE SCALES

The sun delivers ~5,800× human civilization in usable exergy. The biosphere captures only ~0.17% of that — yet runs a planetary-scale harvesting system.

Figure 4. Three power scales, on a logarithmic axis. Industrial civilization delivers ~20 TW of high-exergy power. The biosphere generates ~130 TW of gross chemical-bond formation (most of which cycles back to CO_2). Solar exergy at Earth's surface is roughly 115,000 TW

after albedo, about 5,800× human energy consumption. Energy is not the bottleneck. Harvesting technology is.

The mechanism is precisely what one would build if one set out, with full knowledge of thermodynamics, to engineer a gradient harvester. A photon arrives carrying approximately 2.5 electron-volts of available work. A chlorophyll molecule, structured by roughly 10^4 bits of evolved protein information per reaction center, selectively absorbs the photon and channels its energy into a quantum-coherent excitation. An electron transport chain, again structured by information, ratchets the excitation through a sequence of redox potentials. The terminal output is a chemical bond (ATP, NADPH, glucose) at room temperature, in water, with no flame, no plasma, and no waste byproduct that the planet cannot absorb.

This is a Szilard engine at industrial scale. The hot reservoir is the sun at 5800 K. The cold reservoir is the ambient environment at 300 K. The structured matter is the photosystem. The information embedded in protein and chlorophyll architecture is what makes the harvest selective and efficient. The fact that biology built it from amino acids rather than silicon is incidental. The physics is exactly the physics that any artificial gradient harvester must obey.

The biosphere did not fail to discover combustion

It is sometimes suggested that biology chose gradient harvesting because it could not access mass-destruction physics. This is not what happened. Biology does perform controlled oxidation; cellular respiration burns sugars with oxygen and captures approximately 40 percent of the free energy in ATP. But it does so at 37 degrees Celsius, mediated by enzymes, never as open flame. Biology found a way to extract the free energy of the C-H bond without the destructive high-temperature regime.

Biology did not access fission or fusion. Those regimes require temperatures and conditions that are incompatible with self-assembling, self-repairing molecular machinery. Life systematically avoided the mass-destruction regime not from inability but from architectural incompatibility: the regime cannot coexist with the kind of structure that life is.

This invites an inference, and the inference should be labeled honestly for what it is. The observation that biology avoided mass-disassembly is physics: those regimes are genuinely incompatible with self-assembling molecular machinery. But the step from there to 'there may be a category of technological civilization compatible with biospheric integrity, and a category that is not' is a philosophical extrapolation, not a theorem. It is not derived from the physics; it is suggested by it. Stated as a conjecture rather than a result: mass-disassembly energy, even in its cleanest form, may belong to a class of technologies that are difficult to reconcile with a living planetary system, while gradient harvesting may belong to a class that is naturally compatible with one. This is worth taking seriously precisely because biology ran the experiment for billions of years. But it is an interpretive claim about technology and civilization, and a reader should hold it to the standard of a well-motivated hypothesis, not to the standard of the thermodynamics in the preceding sections.

> **LIFE IS NOT SOLAR-POWERED. LIFE IS GRADIENT-POWERED.**
>
> Photosynthesis is the most visible form of biological harvesting, but it is far from the only one. Chemoautotrophs at hydrothermal vents harvest chemical disequilibria ($H_2 + CO_2$, or $H_2S + O_2$) with no sunlight at all. Electroautotrophs draw electrons directly from minerals. Thermophiles harvest temperature gradients near vents. These organisms populate the deep biosphere, kilometers below the surface, in total darkness, possibly comparable in biomass to all surface life.
>
> The reservoir is therefore much larger than sunlight. The common feature across every form of life is the harvesting of free-energy gradients using informationally-structured matter. Photons are one species of the genus; chemical, thermal, electrical, and osmotic gradients are others. Wherever the

> universe holds a difference that has not yet relaxed, structured matter can be built to harvest it. That is the access point this essay is naming, and it is everywhere.

VI THE STRUCTURAL PARALLEL

Artificial Intelligence is to biological brains what Artificial Energy is to biological photosynthesis. This is not metaphor. It is a structural parallel: a shared logical form between two domains that does not require identical implementations.

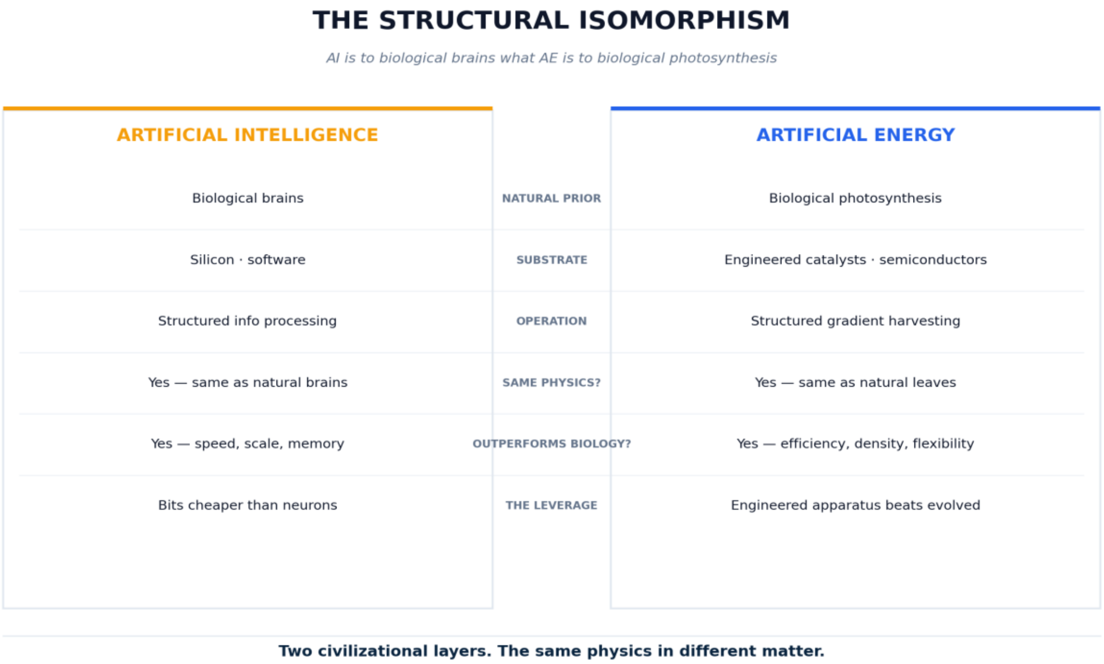

THE STRUCTURAL ISOMORPHISM

AI is to biological brains what AE is to biological photosynthesis

ARTIFICIAL INTELLIGENCE		ARTIFICIAL ENERGY
Biological brains	NATURAL PRIOR	Biological photosynthesis
Silicon · software	SUBSTRATE	Engineered catalysts · semiconductors
Structured info processing	OPERATION	Structured gradient harvesting
Yes — same as natural brains	SAME PHYSICS?	Yes — same as natural leaves
Yes — speed, scale, memory	OUTPERFORMS BIOLOGY?	Yes — efficiency, density, flexibility
Bits cheaper than neurons	THE LEVERAGE	Engineered apparatus beats evolved

Two civilizational layers. The same physics in different matter.

Figure 5. The point-for-point correspondence between AI and AE at the categorical level. Two civilizational technology layers, each beginning with an evolved biological precedent

and proceeding through engineered non-biological substrates. The technological families differ; the logical form is the same.

The parallel between AI and AE is not a slogan, and it is not a perfect isomorphism. It is a structural correspondence between two domains that share a common logical form. Both have an evolved biological precedent. Both are now being engineered in non-biological substrates. Both perform the same underlying physical operation as their biological prior, by the same physics, in different matter. Both have demonstrated, in their early forms, that engineered apparatus can outperform evolved apparatus on specific axes. Both are reaching a moment of civilizational scale.

One asymmetry should be acknowledged immediately. AI's substrate is highly unified: almost all modern AI runs on digital computation in von Neumann or related architectures. AE's substrate is more fragmented. A silicon photovoltaic cell, an engineered enzyme, a heat pump, and a designed catalyst do not share a single underlying substrate the way transformer models and convolutional networks do. The unifying thread of AE is not a common physical implementation but a common operational principle: gradient harvesting in informationally-structured engineered matter. This is parallel to AI at the operational level (structured information processing in engineered substrates) without being identical at the implementation level. The category is real; the technological family is heterogeneous. Both things can be true.

Consider the rows of the correspondence in turn. The natural prior for AI is the biological brain: a hundred billion neurons, evolved over hundreds of millions of years, performing pattern recognition and inference. The natural prior for AE is biological photosynthesis and metabolism: chlorophyll, photosystems, electron transport chains, and enzymes, evolved over billions of years, performing solar-to-chemical energy conversion and selective chemistry. In both cases, the engineered version began with the biological precedent as proof that the operation was physically possible, and then proceeded to find different architectures that could perform the same operation in silicon or designed materials.

The substrate for AI is silicon and software: designed semiconductor architecture running designed algorithms. The substrate for AE is a heterogeneous family of designed materials: silicon photovoltaics, engineered enzymes, designed catalysts, artificial reaction centers, biomimetic membranes. The same kind of design discipline that has driven AI, informationally-structured matter built by human craft to perform a specific computational or thermodynamic function, drives AE.

The operation in AI is structured information processing: taking inputs, applying learned weights, producing outputs that exhibit pattern recognition, inference, language. The operation in AE is structured gradient harvesting: taking ambient free-energy gradients, applying engineered architectures that selectively couple to them, producing useful work or storable chemical bonds. Both operations are governed by the same underlying physics as their biological versions. A silicon neuron and a biological neuron are both implementing structured information processing under uncertainty. A silicon photovoltaic and a chlorophyll molecule are both implementing structured photon harvesting against an entropy gradient.

Most importantly, both AI and AE in their engineered forms can outperform their biological priors on specific axes. AI exceeds biological intelligence on speed, on memory capacity, on parallel scale, on certain narrow pattern-recognition tasks. AE exceeds biological photosynthesis on energy density (silicon PV averages around 200 watts per square meter at peak, against 0.5 watts per square meter for a leaf), on conversion efficiency (47% in laboratory multijunction PV against ~6% for a typical C3 plant), and on output flexibility (electricity, hydrogen, ammonia, designer molecules). The engineered apparatus exceeds the evolved apparatus precisely where evolution faced constraints that human engineering does not face: molecular self-assembly, ambient temperature only, materials limited to organic chemistry.

> ***Two civilizational technology layers, sharing a logical form: a biological precedent, engineered in a non-biological substrate, performing the same operation by the same physics in different matter.***

VII WHERE THINGS STAND NOW

AE is not at a single point on its development curve; its layers are at different stages. The first layer is already scaling. The layers above it sit roughly where AI sat in its pre-takeoff years.

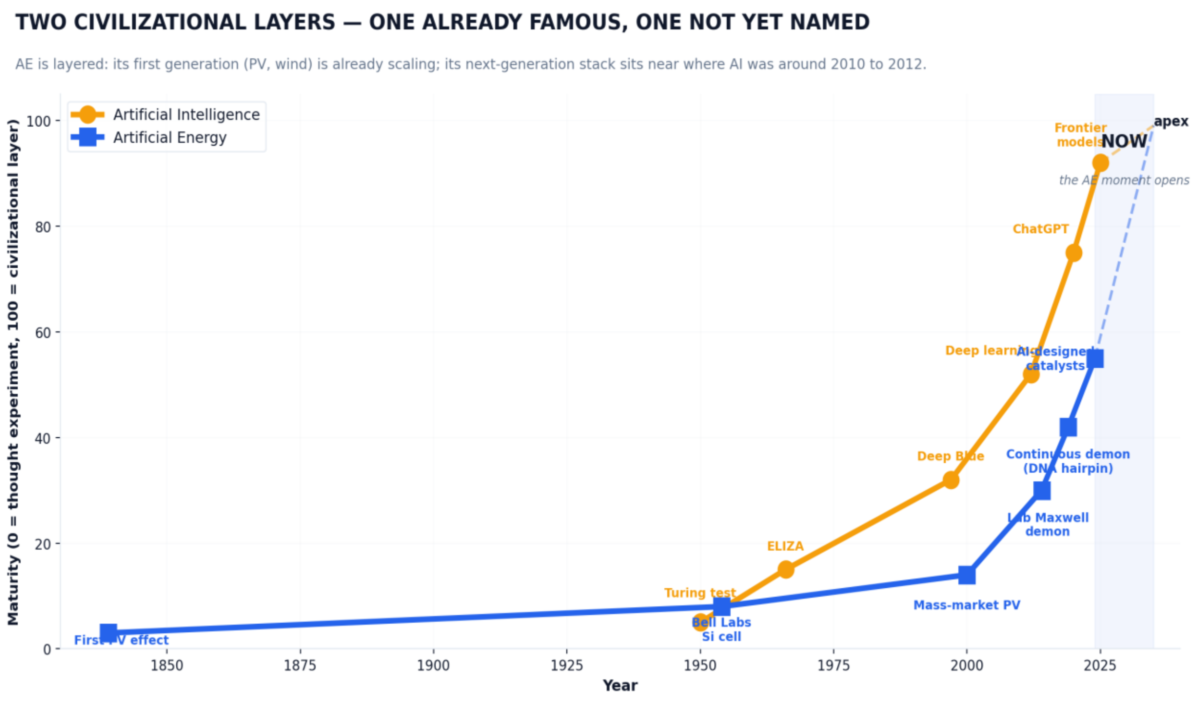

Figure 6. Two civilizational technology curves. AI has had its public moment, beginning around 2012 with deep learning and reaching mass awareness with ChatGPT. The low-information harvesters (photovoltaics, wind) are already scaling, but the information-rich

AE technologies that the thesis is actually about have not yet had their public moment. The technical pieces are emerging: laboratory Maxwell demons, artificial photosynthesis prototypes, AI-designed catalysts. The framing is the missing piece.

A precise statement of where AE stands requires separating two things that are easy to conflate: the regime and its first crude instruments. Photovoltaics and wind belong to the gradient-harvesting regime, and they have already scaled. But they sit at the very floor of it. A bare solar panel is a fixed band gap; a wind turbine is a shaped blade. Each harvests a single gradient in a single way, with almost no information in the apparatus, and each runs quickly into a hard physical ceiling: silicon photovoltaics convert about 22 percent of incoming sunlight at mass-market scale, against a single-junction limit near 34 percent, leaving perhaps a factor of two of headroom and no more. That bounded ceiling is the ceiling of a near-zero-information harvester, and it is the reason the future of this field does not lie in larger panels. Artificial Energy is what happens when the harvesting apparatus becomes informationally rich: when designed catalysts, engineered enzymes, and AI-specified materials begin to select, route, and transform energy flows with the precision that a chlorophyll molecule brings to a photon, and reach gradients that no panel or turbine can touch. Those technologies sit at roughly the equivalent of AI's 2010 to 2012 moment: the underlying capabilities exist in laboratories, the public framing has not yet formed, and the exponential has not yet begun.

Tier 2 — wind and first-generation solar — is not the substance of this argument. These technologies are Passive Gradient Energy: they proved that engineered matter can draw work from an ambient gradient at industrial scale, and they occupy the lowest rung of what gradient harvesting can become. They are likely to be remembered as transitional, the way the steam engine was transitional within the broader electrification of work — necessary, important, and not the destination. The destination is Tier 3: information-rich engineered matter that selects, routes, and transforms gradients the way biology does. That is the ladder whose upper reaches have barely been built.

The Information-Rich AE: What Comes Next

The information-rich forms of AE do not yet exist at industrial scale. They are visible in laboratories and prototypes, in the same way that frontier AI models were visible in research labs years before becoming household terms.

Artificial photosynthesis. Devices that take sunlight, water, and CO_2 as inputs and produce hydrogen, hydrocarbons, or other storable fuels as outputs, at ambient temperature, with no moving parts. Daniel Nocera's group at Harvard has demonstrated solar-to-hydrogen efficiencies of approximately 10 percent. The Joint Center for Artificial Photosynthesis (JCAP), EPFL, Imperial College London, and several private programs are pushing the technology. None has reached industrial scale. This is exactly the kind of underfunded research direction that AE thinking identifies as high-leverage.

Engineered enzymatic catalysis. Biology fixes atmospheric nitrogen into ammonia at room temperature using the nitrogenase enzyme. Industrial civilization fixes nitrogen using the Haber-Bosch process at 400-500 degrees Celsius and 150-300 atmospheres of pressure, consuming approximately 1-2 percent of global energy. The same chemistry performed at radically different thermodynamic costs depending on whether informationally-structured matter is doing the work. AI-designed enzymes that perform Haber-Bosch chemistry at ambient conditions would convert one of the largest energy consumers in industry into an AE process.

AI-designed catalysts. Machine learning systems trained on reaction outcome data can discover catalytic structures that brute-force chemistry cannot. The information embedded in a designed catalyst (its surface geometry, its electronic structure, its coordination environment) directs reactions to outcomes that would otherwise require extreme temperatures or pressures. The catalyst is the apparatus; the AI model is the design tool. This is AI building AE: the most direct manifestation of the parallel.

Multijunction and concentrator photovoltaics. Each junction in a multijunction PV device is tuned to a specific photon energy band, embedding far more structural information than a single-junction panel. This is why multijunction lab devices reach 47 percent efficiency where mass-market silicon reaches 22 percent.

Thermophotovoltaics. Devices that convert any source of heat (concentrated sunlight, geothermal, industrial waste heat) into electricity via tuned photon absorption. MIT demonstrated 40 percent efficiency in 2022. This is gradient harvesting from thermal sources rather than direct solar.

Heat pumps. The clearest industrial-scale example of information directing energy. A heat pump uses information about a temperature gradient to move heat against the gradient, delivering 3 to 5 joules of heat for each joule of electrical work input. A natural-gas furnace delivers approximately 0.95 joules of heat per joule of fuel. The heat pump is 3 to 5 times more thermodynamically efficient because it does not generate heat. It harvests heat that already exists. This is a working, industrial-scale Maxwell's demon already deployed in millions of buildings.

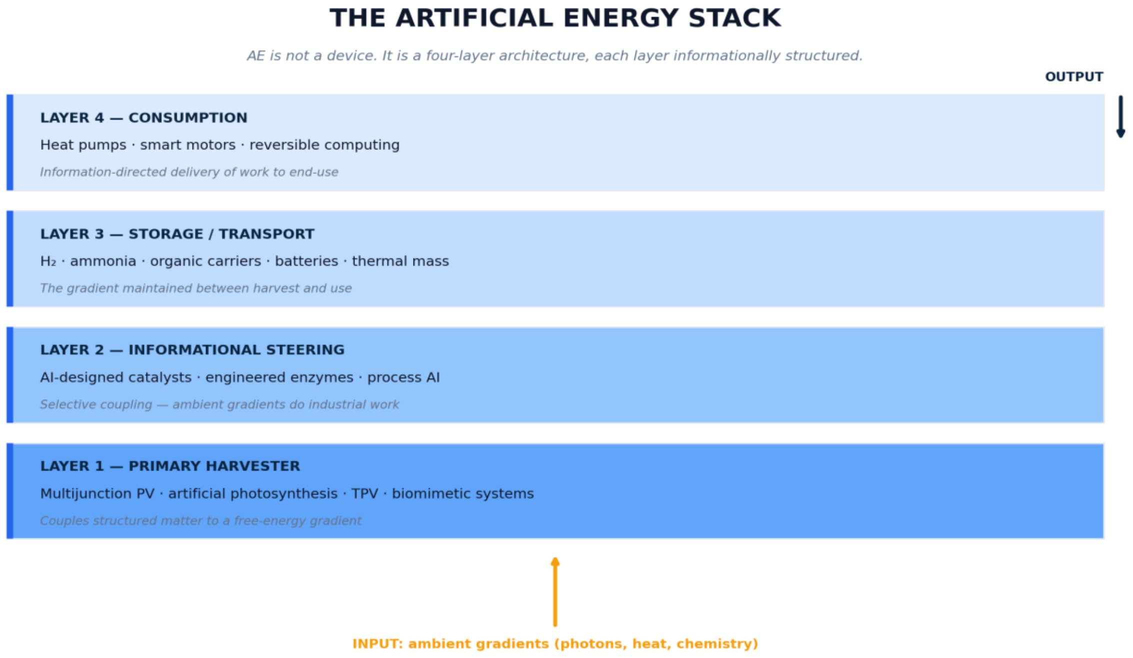

Figure 7. The Artificial Energy stack. AE is not a single device. It is a four-layer architecture in which information is structured into every layer, from primary harvest to end-use consumption.

Two Different Questions About Cost

A first-principles essay should put numbers on its central claims, and it should be careful which claim each number supports. Two distinct cost questions arise in this argument, and they must not be confused. The first is the cost of information used to steer a process. The second is the cost of building and running the physical apparatus that does the harvesting. These are governed by different physics and differ by roughly twenty orders of magnitude, so they require separate treatment.

On the first question, there is a clean physical floor. Steering a chemical process requires handling information, and the minimum thermodynamic cost of irreversibly handling one bit is the Landauer bound, kT ln 2, about 2.87×10^{-21} joules at 300 K. Breaking a chemical bond costs vastly more: the canonical carbon-hydrogen reference is about 6.86×10^{-19} joules. Their ratio is approximately 240. Deciding whether to break a C-H bond is, at the thermodynamic floor, roughly 240 times cheaper than breaking it. This is the bond-bit ratio, derived in full in the canonical source (Anderson, *The Bond-Bit Ratio*, 2026), and it holds for the C-H reference at 300 K; it is not a universal constant across all bonds or temperatures. In real systems the gap widens further, because the relevant information lives in the structure of a designed catalyst or enzyme rather than in transient bit-flips, and a reaction center's design information is paid for once in manufacture and then directs trillions of catalytic cycles. The ratio that matters for AE is the cost of embedded design information to the energy of every reaction it subsequently steers. But this ratio concerns the cost of the decision, not the cost of the machine that executes it. It does not, and cannot, describe the energy required to manufacture a harvester.

The second question . . . the cost of the apparatus . . . is where the regimes diverge most consequentially. For a fuel-consuming plant, the embodied energy of the plant is only the smaller part of the lifetime energy budget; the larger part is the fuel, consumed continuously, every joule delivered requiring a joule or more of fuel burned, forever. For a gradient harvester, the fuel term is zero. The apparatus is paid for once in embodied energy, and the harvest accrues for the operating life of the device against a fuel cost that never comes due. The marginal cost of sunlight, of an ambient temperature gradient, of an unrelaxed chemical disequilibrium, is zero — the sun does not send an invoice, and the gradient is not invoiced either. That asymmetry, between a regime whose fuel must be perpetually consumed and a regime whose fuel arrives whether or not it is captured, is the structural difference the essay rests on. The specific embodied-energy numbers for any particular harvester are an engineering question that depends on which harvester and which manufacturing pathway; the structural point survives any specific number.

VIII WHY “ARTIFICIAL ENERGY” IS THE RIGHT NAME

Naming matters. The wrong name closes a category. The right name opens it.

The conventional vocabulary for energy harvesting has not served the field well. 'Solar' restricts attention to one source. 'Renewable' frames the technology in terms of what it is not (it is not fossil) rather than what it is. 'Clean energy' wraps the technology in moral framing that activates political identities and obscures the underlying physics. None of these terms places the technology in the same cognitive category as Artificial Intelligence: the category of civilization-defining technology layers.

'Artificial Energy' makes that placement explicit. It is truthful: the same word, with the same etymology, applied to the same kind of operation. It is structurally accurate, since AI and AE are closely parallel. It is bold, claiming peer status with the most-discussed technology of the present era. And it is operationally useful, because it tells engineers, investors, policymakers, and the public that the right mental model for energy is no longer 'which fuel to dig up' but rather 'which informationally-structured architecture to design.'

The reasons follow, in order of strength.

Reason 1: It Is True in the Same Sense AI Is True

Artificial Intelligence is intelligence engineered in non-biological substrates. Artificial Energy is gradient harvesting engineered in non-biological substrates. The etymology of 'artificial' is the same in both cases: Latin artificialis, made by skill. The structure of the analogy is identical. The physics is identical. If AI is the right name for one, AE is the right name for the other. To accept one and reject the other requires a special argument that no one has yet supplied.

Reason 2: It De-Categorizes the Topic from 'Fuel Choice'

When the public hears 'solar energy,' they place it in a mental category labeled 'energy sources,' one option in a list including coal, gas, nuclear, wind. It is a fuel choice. It sits next to coal in the cognitive map—and solar panels, correctly understood, are Tier 2: passive gradient harvesters that are better than Tier 1 but are not the civilizational leap the energy transition requires. When the public hears 'Artificial Intelligence,' they do not place it in a category of 'kinds of calculation.' They place it in a category called 'civilizational technology layers,' alongside electricity, the internet, the printing press. 'Artificial Energy' forces the same cognitive move. AE is not a fuel choice. It is not a better solar panel. It is a domain, peer to AI, operating at Tier 3: the information-rich frontier that solar and wind cannot reach.

Reason 3: It Rides the AI Mindshare

Approximately half a trillion dollars of capital and an enormous concentration of generational talent are flowing into AI. The infrastructure that has been built for AI (compute clusters, design tools, training pipelines) is largely the same infrastructure needed to design AE: machine learning systems for materials discovery, simulation tools for molecular dynamics, design optimization for engineered enzymes and catalysts. Naming AE as the natural sibling to AI is not just rhetorical. It is operationally efficient. The same teams, the same companies, the same capital pools, the same playbook of 'engineer a non-biological version of a biological capability and scale it' can apply directly. The intellectual current is already running. AE harnesses it.

Reason 4: It Is Honest About the Limits

'Artificial Energy' is also honest about what AE is not. It is not a new fuel. It is not energy from nothing. It is not a violation of conservation laws. Just as AI processes real, pre-existing information through engineered architectures, AE harvests real, pre-existing energy gradients through engineered architectures. The substrate is artificial. The phenomenon is real. The 'artificial' word is a flag for the origin of the apparatus, not a claim that the energy is fake or free. This honesty protects the framing from criticism: no competent physicist will reject 'Artificial Energy' as a violation of the First Law, because the term does not claim to violate it. The energy comes from gradients. The apparatus that harvests it is artificial. That is all.

Reason 5: It Activates the Right Emotion

There is no point pretending that civilizational shifts happen on physics alone. They happen when the framing creates urgency, alignment, and excitement. AI has the urgency, alignment, and excitement. AE, properly named, inherits all three. Engineers will want to build it. Investors will want to fund it. Governments will want to support it. Students will want to enter the field. The fear of missing the next big thing, the same fear that has driven every dollar into AI since 2022, will, properly directed, drive the deployment of AE at the speed civilization needs.

Reason 6: It Is the Simplest Name That Captures the Regime

Among the alternatives ('negentropic harvesting,' 'photonic energy,' 'gradient energy,' 'information-mediated energy,' 'biomimetic energy'), only 'Artificial Energy' captures the regime in two words, anchors to a familiar cognitive category, and remains technically accurate. Simplicity is not a cosmetic feature. It is a functional feature. Names that do not fit in headlines do not propagate.

> **THE DEFENSIBLE CLAIM, STATED PRECISELY**
>
> Artificial Energy is energy harvest engineered in non-biological substrates, doing what biological photosynthesis and metabolism do (selective coupling of structured matter to existing free-energy gradients) by the same physics, in different matter. It is not a new energy source. It does not violate any conservation law. It is the engineered version of biology's gradient-harvesting strategy, and it is the natural peer of Artificial Intelligence as a civilizational technology layer.

IX THE CLAIM THAT MUST NOT BE MADE

The First Law of Thermodynamics has no exceptions. The case for Artificial Energy does not require any.

It is tempting to formulate the argument more aggressively than the physics allows. One might be tempted to say that information is itself a new energy source, that civilization can be powered by bits, that AE makes energy out of information the way AI makes intelligence out of computation. None of these statements is true.

Information has no intrinsic energy. A bit is a logical state of some physical degree of freedom; the physical degree of freedom has thermal energy, but the logical state itself has no energy. A '0' and a '1' weigh the same. They store the same energy. The First Law forbids creation of energy from nothing, and no rearrangement of information violates this. The Toyabe-Koski-Ribezzi experiments did not produce energy from information; they used information to harvest energy that was already present in a heat bath, with full thermodynamic bookkeeping balancing every joule.

Per quantum carrier, the cost of irreversibly handling a bit at 300 K is approximately 18 millielectron-volts: about 140 times less than a single visible photon, and roughly 240 times less than the canonical carbon-hydrogen bond, the bond-bit ratio established earlier. That ratio is specific to the C-H reference at 300 K, not a universal constant; weaker bonds narrow it and the Landauer cost scales with temperature. The general point survives the qualification. What makes information valuable is not how much energy it carries, since it carries none, but how cheaply it can be handled relative to the physical changes it can direct. A small, well-placed quantity of information can steer the harvest of much larger energies, at a handling cost far below the energies in play.

The correct claim is this. Artificial Energy is the engineered version of biology's gradient-harvesting strategy. It uses information as the design and steering technology that makes the harvest selective and efficient. It does not replace the need for free-energy gradients. It deploys structured matter to harvest gradients that already exist. This is precisely what photosynthesis does. It is precisely what photovoltaics do. It is precisely what every member of the AE family does and will do. The substrate is artificial. The gradients are natural. The energy is real.

> ***Information is not a new energy source. Information is the technology that turns existing gradients into useful work without destroying matter. The biosphere has run on this principle for 3.5 billion years. Industrial civilization has not yet noticed.***

X WHAT THIS MEANS

The argument is not 'stop building reactors tomorrow.' The argument is that civilization is overinvested in the primitive regime and underinvested in the mature one.

The mass-destruction regime remains necessary as a transitional bridge. Photovoltaic deployment requires steel, which requires combustion-derived process heat. Grid storage at planetary scale is not yet solved. Cement, ammonia, and certain industrial chemistry have not yet been electrified at scale. Deep-space propulsion, where solar flux falls as the inverse square of distance, may always require mass-based fuels. Fission remains a legitimate decarbonization tool during the transition. Fusion, if it can be made to work, will be an important long-term backup for niches that ambient gradients cannot serve.

But these are bridge applications and niche applications. They are not the destination. The destination is the regime that has already been validated by 3.5 billion years of biological evolution at 130 terawatts of continuous throughput: gradient harvesting via informationally-structured matter, deployed deliberately by humans for the first time in history.

The capital allocation argument follows directly, though it should be stated at the right level of confidence. Fusion R&D currently receives approximately $5-10 billion per year globally in public funding, plus a comparable amount from private companies. Advanced fission and fission research receive another few billion. Combustion-related research (better gas turbines, carbon capture, methane leak reduction) receives tens of billions. Artificial photosynthesis receives on the order of $100 million per year globally. Engineered enzymatic catalysis is dispersed across many programs but underfunded relative to its potential leverage. Reversible computing, which would allow computation to approach the Landauer floor and unlock much greater AE multipliers in the consumption layer, is a tiny research program.

There is a real possibility that some of these underfunded fields are underfunded because the physics is genuinely hard. Artificial photosynthesis has been described as 'ten years away' since the 1970s; that history should be respected, not dismissed. The argument here is not that AE technologies are obviously high-leverage. It is hypothesis-level: the underfunding-versus-potential pattern in AE today resembles the pattern in AI during the 1990s and 2000s, when decades of slow basic research suddenly compounded once the underlying capabilities crossed a threshold. The pattern is not a guarantee. But the asymmetry between dollars flowing into mass-disassembly research and dollars flowing into the layered AE stack is large enough that it is worth examining seriously, not because anyone knows which AE bets will pay off but because the prior allocation appears poorly calibrated to where physical leverage exists.

The leverage is plausibly on the AE side. The money is overwhelmingly on the mass-disassembly side. This may be a misallocation. The right response is not to defund the mass-disassembly path, which retains legitimate transitional uses, but to test the AE hypothesis with capital commensurate with its potential, the way AI was eventually tested when sustained patient investment in deep learning began to compound.

Two civilizational technology layers are being built in this generation. One processes information. The other harvests energy. Both are engineered versions of what biology has been doing for billions of years. AI is the first to be widely named and widely capitalized. AE is the second. The framing matters. So does the name.

XI TRADITIONAL ENERGY VS ARTIFICIAL ENERGY

The two regimes can be compared directly, on numbers that hold. Some of the comparison is fixed by physical law; some is empirical and current. Both are stated here, with the difference between them marked.

Throughout this essay, traditional energy means the substrate-consuming regime: combustion of coal, oil, and gas, and the fission of heavy nuclei. Artificial Energy means the substrate-preserving regime: engineered gradient harvesting. The contrast below runs across four measurable dimensions. The first two are governed by physical law and are not matters of opinion or forecast. The second two are empirical, drawn from current measured data, and could shift as technology and markets move; they are flagged accordingly in the notes.

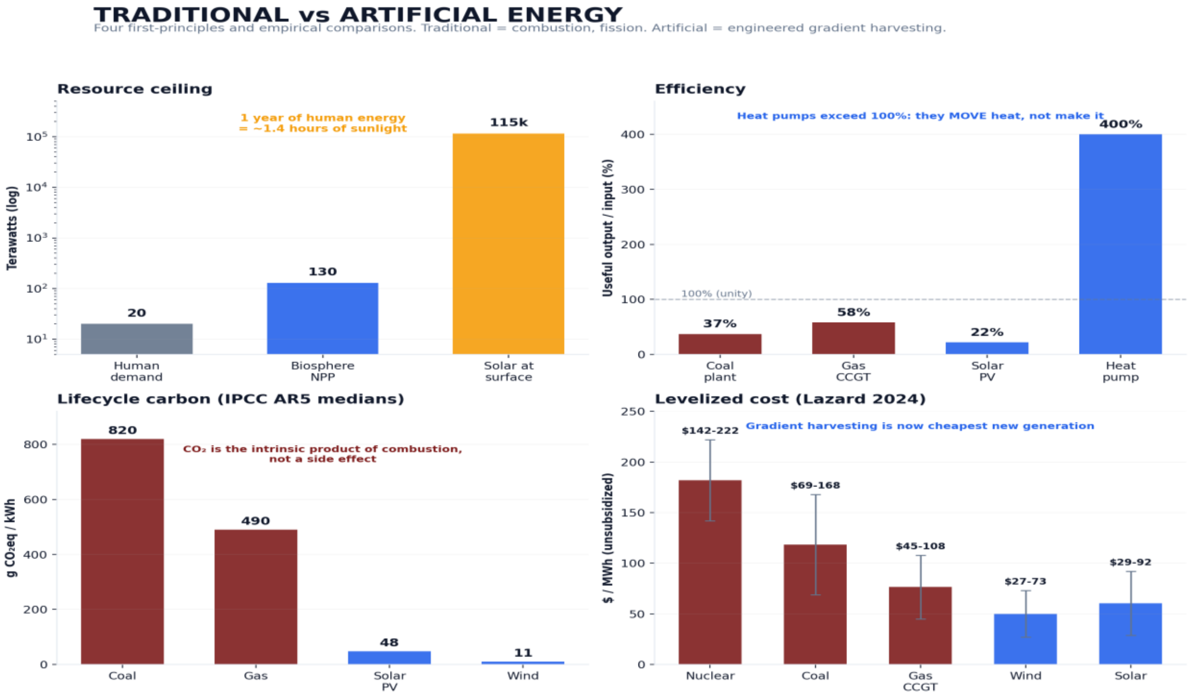

Figure 8. Four comparisons between the regimes. Resource ceiling and efficiency are fixed by physics; lifecycle carbon and levelized cost are empirical (IPCC AR5 medians and Lazard 2024, respectively). Red marks traditional energy; blue marks Artificial Energy.

DIMENSION	TRADITIONAL ENERGY	ARTIFICIAL ENERGY
Resource ceiling	Finite stock; depletes as used	~115,000 TW solar at surface; renewed daily
Energy return (EROI)	Declines as easy reserves deplete	PV ~10-30, wind ~15-35, and rising
Efficiency limit	Carnot-bound; must dump >40% as heat	Not Carnot-bound; heat pumps reach 300-500%
Marginal fuel cost	Floored by extraction; never zero	Zero; the gradient is not invoiced
Lifecycle carbon	Coal ~820, gas ~490 g/kWh	Solar ~48, wind ~11 g/kWh
Levelized cost (generation)	Coal $69-168, gas $45-108 /MWh	Solar $29-92, wind $27-73 /MWh
Waste profile	CO2 intrinsic to the reaction	Recycled material flows; heat re-radiated

What the Physics Fixes

The first two rows are not negotiable. The resource ceiling is a flux-versus-stock distinction: traditional energy draws on a finite inventory of reduced carbon and heavy nuclei that took hundreds of millions of years to accumulate and is consumed on a timescale of centuries, while gradient harvesting draws on an incident flux that is renewed continuously. The sun delivers roughly 115,000 terawatts of usable exergy to Earth's surface against a human demand near 20 terawatts. One year of total human energy consumption equals about one and a half hours of sunlight falling on the planet. The entire remaining recoverable stock of fossil fuel, on the order of 10^{22} joules, is matched by a few days of incident sunlight. This is not a claim about better technology. It is arithmetic about flux and stock.

The efficiency ceiling is equally fixed, and it is where the regimes diverge most sharply at the level of principle. A heat engine that produces work by burning fuel is bound by the Carnot limit: it must reject a large fraction of its input as waste heat, no matter how well engineered, because the second law requires it. A coal plant converting 37 percent of its fuel energy to electricity is already near the practical ceiling for its operating temperatures. Gradient harvesting is not bound by the same limit, because it is not running a heat engine between two reservoirs in the same way. The clearest demonstration is the heat pump, which delivers three to five units of heat for every unit of work it consumes, a coefficient of performance of 300 to 500 percent. This does not violate conservation of energy; the additional heat is moved, not created, drawn from the ambient environment by a device that uses information about a temperature difference to pump against it. No combustion process can exceed unity in this sense. Many harvesting processes routinely do.

What the Data Shows

The second two rows are empirical rather than law-bound, and they are stated at the appropriate confidence. On lifecycle carbon, the asymmetry is structural even though the numbers are measured rather than derived: carbon dioxide is the intrinsic product of combustion, the unavoidable output of the reaction itself, whereas the lifecycle emissions of gradient harvesting come almost entirely from manufacturing and fall as the manufacturing grid decarbonizes. The Intergovernmental Panel on Climate Change's median lifecycle figures are about 820 grams of CO2-equivalent per kilowatt-hour for coal and 490 for natural gas, against about 48 for utility solar and 11 for onshore wind. The gap is between one and almost two orders of magnitude, and on the harvesting side it trends toward its manufacturing floor rather than being fixed by the chemistry of a flame.

The Asymmetry Beneath the Table

Each row of the comparison is a surface expression of a single underlying difference, the one established at the start of this essay. In traditional energy, the fuel and the apparatus are the same matter, so the energy obtainable per unit of fuel is fixed, the stock depletes, and the cost per joule has a hard floor set by extraction. In Artificial Energy, the apparatus and the gradient are separate, so the apparatus persists, the energy it directs accumulates with its operating lifetime, and the marginal cost of the gradient is zero because the sun does not send an invoice. The resource ceiling, the efficiency behavior, the waste profile, and the cost trajectory are not four independent advantages. They are four consequences of building energy systems that harvest rather than consume. That is the difference the table measures, and it is a difference of structure, not of degree.

XII ON THE LIMITS OF THIS ARGUMENT

A serious thesis deserves a serious accounting of where it strains. This section enumerates the weaknesses the careful reader will find on their own.

The taxonomy of mass-disassembly versus gradient harvesting is rhetorically clean, but it imposes a sharper line than the physics strictly draws. Combustion converts roughly four parts in ten billion of its fuel mass to energy; calling that 'mass destruction' alongside fission and fusion, which convert measurable fractions, is taxonomically loose. The distinction that actually carries the argument is whether the fuel and the apparatus are the same matter, whether the energy-producing reaction consumes its own substrate. By that test, combustion, fission, and fusion belong on one side and gradient harvesting on the other. Readers who object to 'mass destruction' as a phrase may substitute 'fuel-consuming' or 'substrate-sacrificing' without changing the argument.

The AI–AE correspondence is a structural parallel, not a perfect isomorphism. AI's substrate is unified: almost all of it runs on digital computation. AE's substrate is heterogeneous, since photovoltaics, engineered enzymes, designed catalysts, heat pumps, and biomimetic systems do not share a single physical implementation. What they share is an operational principle: gradient harvesting in informationally-structured engineered matter. The parallel holds at the level of category, not at the level of technological family. This asymmetry should be acknowledged openly; the AI parallel remains useful as a frame even where it does not survive at every level of detail.

Comparing the biosphere's 130 terawatts of net primary productivity to industrial civilization's 20 terawatts of delivered power overstates the contrast. The two numbers measure different things. Biospheric NPP is gross chemical-bond formation, much of which decays back to CO_2 on short timescales. Industrial energy is delivered, dispatchable, high-exergy power. The honest comparison is between human demand (~20 TW) and total solar exergy at Earth's surface (~115,000 TW), a ratio of roughly 5,800. The biosphere's role in the argument is as existence proof, showing that gradient harvesting works at planetary scale, not as a demonstration that biology produces more useful power than industry does.

The capital-allocation claim is hypothesis-level, not the conclusion of an expected-value calculation. Underfunded fields can be underfunded because the physics is genuinely hard. Artificial photosynthesis has been ten years away since the 1970s. The claim here is that the underfunding pattern in AE resembles the underfunding pattern in AI during the long pre-takeoff decades, not that AE is guaranteed to compound on the same trajectory. That guarantee would require evidence that does not yet exist.

Finally, the very strength of the AE framing, the AI parallel, is also a constraint. If a higher-platform figure introduces a different name for the same regime, that name will win regardless of priority. 'Engineered photosynthesis,' 'post-combustion energy,' or another framing could displace 'Artificial Energy' for reasons unrelated to which is more accurate. The argument here does not depend on the name surviving. It depends on the regime being recognized as a distinct civilizational layer, by whatever name.

The most serious physical objection is power density, and the argument must concede it plainly. Mass-disassembly's genuine advantage is concentration. A kilogram of uranium or a tank of jet fuel carries enormous energy in a small volume and delivers it on demand, which is why combustion and fission still dominate aviation, shipping, heavy industry, and dispatchable baseload power. Gradient harvesting is diffuse: sunlight arrives at roughly a kilowatt per square meter at noon and far less on average, wind carries only a few watts per square meter of land, and capturing usable power at scale therefore demands large areas, large material flows, and critical minerals. The flux is abundant in total, as the resource-ceiling figures show, but it is thin per unit area, and thinness is a real engineering cost that the abundance figures alone do not capture.

Intermittency compounds this. The sun sets and the wind drops, while demand does not, so a harvesting-dominated system must store energy or hold dispatchable backup to deliver power on demand. Grid-scale storage carries its own embodied energy, material, and round-trip losses, and the levelized-cost figures cited earlier describe energy at the point of generation, before these firming costs are added. Once storage and backup sufficient for reliable supply are included, the cost advantage of harvesting narrows, and in some reliability-matched comparisons the regimes converge. None of this overturns the structural argument, but it bounds how fast and how completely the harvesting regime can displace the consuming one, and it explains why the transition is an engineering project of decades rather than an immediate substitution.

A final limit is one of category rather than physics. This essay makes three different kinds of claim, and they carry different weights. The physical taxonomy, that some systems consume their substrate while others harvest a persistent gradient, is the strongest and rests directly on thermodynamics. The technology forecast, that engineered gradient harvesting will mature into a major civilizational layer, is a reasonable extrapolation from current trends but is not entailed by the physics. The investment prescription, that civilization is overinvested in the consuming regime, is advocacy: a judgment about what ought to follow from the first two claims. The normative vocabulary this essay uses, calling one regime 'mature' and the other 'primitive,' naming a 'destination,' describing civilization as 'overinvested,' belongs to that third category. These are arguments the author believes are well supported, but they are strategic claims stated in the cadence of physical ones, and a careful reader should weigh them as such rather than treating them as consequences of Landauer's bound or the solar flux.

None of these limits is fatal. Taken together, they bound the strength of the claim without undermining it. The strongest and most defensible version of the thesis is this: engineered systems can harvest the free-energy gradients that the universe continuously supplies, using persistent, information-rich structures, biology proves this can operate at planetary scale, and the regime has a fundamentally different and more favorable cost structure than substrate consumption because its fuel term is zero. The weakest version, that civilization has discovered a new thermodynamic category whose dominance follows automatically from information theory, is not established and is not claimed here. What is claimed is that the regime is real, that its engineered forms are underdeveloped relative to their potential, and that the enterprise deserves a name placing it in the category of civilizational technology layers rather than fuel choices.

XIII WHAT HAS ACTUALLY BEEN FOUND

It is worth stating plainly what is, and is not, being claimed as a discovery, because the claim is easy to mistake in both directions.

Nothing in this essay is new physics. There is no new law, no new mechanism, no violation of any conservation principle, no escape from the second law. Photosynthesis, catalysis, the Landauer bound, photovoltaics, heat pumps, and the information-thermodynamics experiments are all established science. A reader looking for a breakthrough in the laws of nature will not find one here, and none is claimed.

What is being claimed is a discovery of a different kind, the kind that reorganizes scattered facts into a single thing that was always present but never named. There is an access mode to energy that civilization has been using in fragments, under a dozen different headings, without recognizing that the fragments are one thing. Artificial photosynthesis, designed catalysts, engineered enzymes, artificial reaction centers, biomimetic membranes, adaptive materials, the whole emerging family of information-rich gradient harvesters, are not a miscellany of clean-energy projects. They are early instances of a single regime: the deliberate engineering of matter, structured with enough information, to couple selectively to the free-energy gradients that already pervade the universe and to draw work from them without consuming itself. That regime is the oldest source of organized energy on Earth. Life has run on it for three and a half billion years. What is new is not the source. What is new is the recognition that humans can now reach it on purpose, by design, at industrial scale, and that doing so is one coherent enterprise rather than a scattering of unrelated efforts.

Seen this way, three things about the access mode follow, and they are the reason it matters. The first is that it is abundant in a way fuel never was. The gradients are everywhere and they are renewed: the sun alone delivers thousands of times human demand, and sunlight is only the largest of many gradients, alongside chemical, thermal, osmotic, and kinetic differences that structured matter can be built to harvest. The binding constraint on this mode has never been the availability of the resource. It has been the sophistication of the apparatus, which is to say the information.

The second is that, in operation, the mode is structurally cheaper and cleaner than substrate consumption, for reasons that are physical rather than incidental. It is not Carnot-limited in the same way a heat engine is; the fuel term in its lifetime energy budget is zero; and it produces no combustion products at the point of harvest because there is no combustion. These are not efficiency tweaks won by better engineering. They are consequences of harvesting a gradient rather than consuming a substrate. The apparatus is built once and then draws, for decades, on a flux that arrives whether or not it is captured.

The third is that the mode is non-destructive in a specific and limited sense that should not be overstated. Building the apparatus costs embodied energy, materials, and mining, and deploying it at scale demands land, storage, and critical minerals, as the limits of this argument make clear. The claim is not zero footprint. The claim is that, unlike combustion and fission, the mode does not require the perpetual destruction of a fuel stock to keep running. The destruction, where it exists, is in the one-time construction, not in the operation. A combustion economy must keep burning forever; a harvesting economy builds and then collects.

This is the find. Not a new law, but a new name for an old and enormous thing that had been hiding in plain sight, broken into pieces too small to see whole. The transformative possibility, energy that is abundant, structurally cheap, and non-consuming, is a genuine possibility and not a promise; the limits of power density, intermittency, and embodied cost stand between the possibility and its realization, and they are real. But a possibility cannot be pursued deliberately until it has been seen as one thing and given one name. That is what this essay has tried to do. The reservoir is older than life. The recognition that it can be reached on purpose, and that reaching it is a single enterprise worth the focus that has gone to drilling and splitting and burning, is the discovery on offer here.

XIV THE SHORT VERSION

ARTIFICIAL ENERGY IN TWO PAGES

For 200 years, we drew most of our energy from disassembling matter. We burned coal. We burned oil. We split atoms. We tried to fuse them. In combustion, fission, and fusion alike, the fuel and the apparatus were always the same matter. The fuel was always consumed. Life runs by a different method, and has for 3.5 billion years.

The biosphere does not consume its apparatus. It uses informationally-structured matter (chlorophyll, enzymes, evolved molecular machines) to harvest gradients that already exist: sunlight, heat, chemistry. At ambient temperature. With no combustion. Recycling its materials rather than accumulating waste. At 130 terawatts of continuous net primary productivity, planetary in scope.

Energy comes in three tiers. Tier 1 is mass-destruction: combustion, fission, fusion—the substrate is consumed, waste accumulates. Tier 2 is passive gradient energy: solar panels, wind turbines, hydro—the apparatus is passive, selects nothing, and hits a hard physical ceiling. Tier 3 is what biology has been doing for 3.5 billion years, and what industrial civilization has not yet done deliberately: information-rich gradient harvesting, where engineered matter selects, routes, and transforms energy at the molecular scale. We are now beginning to engineer the substrate-preserving form deliberately for the first time. Its name is Artificial Energy.

Artificial Intelligence is intelligence engineered in non-biological substrates (silicon, software) doing what biological brains do by the same physics in different matter. Artificial Energy is gradient harvesting engineered in non-biological substrates (designed materials, engineered catalysts, artificial reaction centers) doing what biological photosynthesis does by the same physics in different matter.

This is not a better way to capture sunlight. It is access to a deeper reservoir: the free-energy gradients that fill the universe, harvested by matter structured with enough information to steer them. They are the next civilizational technology layer, parallel to AI in significance.

Stop overinvesting in mass-destruction. Stop mistaking passive solar and wind for the destination. Start building Artificial Energy—the third tier, the one that barely exists yet, and the one that matters.

JED ANDERSON · 2026

REFERENCES

Anderson, J. (2026). The Bond-Bit Ratio: A derivation of why information is at least 240× cheaper than force. https://jedanderson.org/essays/bond-bit-ratio

Bennett, C. H. (1973). Logical reversibility of computation. IBM Journal of Research and Development, 17(6), 525-532.

Bennett, C. H. (1982). The thermodynamics of computation: a review. International Journal of Theoretical Physics, 21(12), 905-940.

Bérut, A., Arakelyan, A., Petrosyan, A., Ciliberto, S., Dillenschneider, R., & Lutz, E. (2012). Experimental verification of Landauer's principle linking information and thermodynamics. Nature, 483, 187-189.

Fraunhofer ISE (2021). Photovoltaics Report. Fraunhofer Institute for Solar Energy Systems. Energy payback times and life-cycle energy analysis for crystalline-silicon PV.

IPCC (2014). Climate Change 2014: Mitigation of Climate Change. Working Group III, Annex III: Technology-specific cost and performance parameters (lifecycle GHG medians).

Hazen, R. M., Burns, P. C., Cleaves, H. J. II, Downs, R. T., Krivovichev, S. V., Wong, M. L. (2023). On the roles of function and selection in evolving systems. PNAS, 120(43), e2310223120.

Jaynes, E. T. (1957). Information theory and statistical mechanics. Physical Review, 106(4), 620-630.

Koski, J. V., Maisi, V. F., Pekola, J. P., & Averin, D. V. (2014). Experimental realization of a Szilard engine with a single electron. PNAS, 111(38), 13786-13789.

Landauer, R. (1961). Irreversibility and heat generation in the computing process. IBM Journal of Research and Development, 5(3), 183-191.

Lazard (2024). Levelized Cost of Energy+ (LCOE+), version 17.0. Unsubsidized levelized cost ranges for generation technologies.

Maxwell, J. C. (1867). Letter to P. G. Tait describing the demon, published in Knott's Life and Scientific Work of P. G. Tait (1911).

McCarthy, J., Minsky, M., Rochester, N., Shannon, C. (1955). A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence.

Nocera, D. G., et al. (2011). Wireless solar water splitting using silicon-based semiconductors and Earth-abundant catalysts. Science, 334(6056), 645-648.

Ribezzi-Crivellari, M., & Ritort, F. (2019). Large work extraction and the Landauer limit in a continuous Maxwell demon. Nature Physics, 15, 660-664.

Schrödinger, E. (1944). What is Life? Cambridge University Press.

Shannon, C. E. (1948). A mathematical theory of communication. Bell System Technical Journal, 27, 379-423, 623-656.

Shockley, W., & Queisser, H. J. (1961). Detailed balance limit of efficiency of p-n junction solar cells. Journal of Applied Physics, 32(3), 510-519.

Szilard, L. (1929). On the decrease of entropy in a thermodynamic system by the intervention of intelligent beings. Zeitschrift für Physik, 53, 840-856.

Toyabe, S., Sagawa, T., Ueda, M., Muneyuki, E., & Sano, M. (2010). Experimental demonstration of information-to-energy conversion and validation of the generalized Jarzynski equality. Nature Physics, 6, 988-992.

Wheeler, J. A. (1990). Information, physics, quantum: The search for links. In Complexity, Entropy, and the Physics of Information. Addison-Wesley.

JED ANDERSON
Writer and builder. Working at the intersection of physics, information theory, and the protection of nature.
jedanderson.org · 2026

A NOTE ON METHOD

This essay was developed with the assistance of artificial intelligence, used as a tool for derivation, verification, drafting, and critique. Every quantitative claim was independently checked; the argument, the judgments, and the conclusions are the author's own. That AI served as an instrument in building the case for Artificial Energy is fitting: the same class of *tool that designs the catalysts and harvesting systems described here also helped articulate why they matter.*