

The Missing Half

Every Vision of AI, Every Documented Harm, and the One Proposal That Builds a Right Hemisphere

J. Konstapel, Constable Research, Leiden · 22 October 2026 · **version 2.0**

© 2026 J. Konstapel, Constable Research, Leiden. All rights reserved.
Quotation with attribution is permitted; reproduction of the whole requires written permission.

Version 2.0 replaces “One Path, Many Worries” (version 1.0).

Abstract

This essay brings together the visions of AI held by 136 named people and institutions. Labs, safety researchers, sceptics, economists, unions, churches, governments and the Global South are all included. It also catalogues 111 documented harms, each with names, places and numbers. Independent checking confirmed every figure that could be traced. No checked claim was refuted.

The visions fall into six groups: accelerate, steer, stop, doubt, change the goal, and secure sovereignty. They disagree about speed, safety and ownership. On one point they agree without saying so. They all assume the same carrier: digital, clocked, binary computation measured in FLOP, gigawatts and capital expenditure. They also all see AI as a single half. It must become larger, smaller, safer or fairer, but it remains one half.

One proposal starts elsewhere. Right-Brain AI calls today’s language models a “hypertrophied left hemisphere”. It proposes the missing right hemisphere as a resonant field of coupled oscillators, computing in phase rather than in bits, on photonic silicon nitride with three phase values.

Four independent searches looked for anyone else who combines these three elements. Together they ran close to seventy queries in five languages and in patents. They found no one. Some groups hold two of the three. Nobody holds all three.

1. The question

In September 2026 a frontier lab wrote into its draft prospectus that advanced AI may pose “catastrophic or existential risks to humanity”. Investors meeting in Amsterdam a few weeks later had no framework for that sentence.

The sentence invites a wider question. Who else worries about AI, and what do they want? Which harms are documented, with names attached? And does any vision step outside the frame they all share?

To answer it, the visions and harms were gathered without a preferred outcome. Section 9 says how. The visions are then bundled by what each party wants. The harms are bundled by who is hit. Finally, both are held up against one proposal that starts from a different picture of what AI is.

The picture used throughout is the fishnet under tension. Every party is a knot. The tension runs along the strands between them. What holds is what closes.

2. Six groups, one carrier

Each person or institution was placed in one group, according to what they want. The checking judges counted 111 positions in the main lines:

- 21 want to accelerate;
- 35 want to steer;
- 11 want to stop;
- 22 doubt the path;
- 15 want a different goal;
- 7 fit nowhere.

A further 25 voices added later were counted separately. In that set, sovereignty appeared as a group of its own.

2.1 Accelerate

This group wants more scale, more energy and more computing power.

- **Sam Altman** (OpenAI): “We are past the event horizon; the takeoff has started.” His wish: “May we scale smoothly, exponentially and uneventfully through superintelligence.” He expects “the cost of intelligence should eventually converge to near the cost of electricity”.

- **Jensen Huang** (Nvidia): “AI is infrastructure.” Also: “Trillions of dollars of infrastructure still need to be built.”
- **Elon Musk** (xAI, now part of SpaceX): “Long term space-based AI is obviously the only way to scale.” He speaks of launching “a terawatt per year”.
- **Sundar Pichai** (Alphabet) also wants data centres in space, “so that we can better harness the energy from the sun”.
- **Andy Jassy** (Amazon): “We’re not investing approximately \$200 billion in capex in 2026 on a hunch.”
- **Mark Zuckerberg** (Meta) wants “personal superintelligence to everyone”, worn as glasses.
- **Masayoshi Son** (SoftBank): “They will be 10,000 times smarter than us.”
- **Ray Kurzweil**: “It won’t be us versus AI. We’re going to be made much more intelligent by merging with AI.”
- **Leopold Aschenbrenner**: “AGI by 2027 is strikingly plausible.” He wants a state project.
- **Marc Andreessen**: “any deceleration of AI will cost lives”.
- **Peter Thiel** calls the opponent of science a “luddite who wants to stop all science”.
- **JD Vance** (Paris, February 2025): “The AI future is not going to be won by hand-wringing about safety.” The US AI Action Plan of July 2025 speaks of “a race to achieve global dominance”.
- **Liang Wenfeng** (DeepSeek): “Our ultimate goal is still to achieve AGI.” Alibaba’s **Eddie Wu**: “Tokens are the electricity of the AI world.”
- **Kai-Fu Lee** puts the gap between the United States and China at about six months.

2.2 Steer

This group wants to keep building, with rules, testing and ethics.

- **Dario Amodei** (Anthropic) sees “a country of geniuses in a datacenter”. He wants transparency laws, export controls and interpretability. “Humanity needs to wake up.”
- **Demis Hassabis** (Google DeepMind) calls the chance that scaling alone suffices “50/50”. He wants “an ‘International CERN for AI’”.
- **Mustafa Suleyman** (Microsoft AI) wants “a subordinate, controllable AI”. “Humans matter more than AI.”
- **Satya Nadella** wants “bicycles for the mind”. He adds that AI needs “societal permission”.
- **Yoshua Bengio** chairs the International AI Safety Report. He builds a non-agentic “Scientist AI” as a check on agentic systems.
- **Geoffrey Hinton** calls a 10% chance of human extinction within ten years “not an unreasonable estimate” (September 2026). He wants AI with a “maternal instinct”.
- **The safety institutes.** The UK AI Security Institute: “We’ve found universal jailbreaks for every system we’ve tested.” METR wants evaluators embedded inside labs. Apollo Research: “The warning shots have been fired.”
- **The IDAIS dialogues** bring together Andrew Yao, Xue Lan, Zhang Ya-Qin, Bengio and Russell. London, April 2026: “Currently deployed AI safeguards are far from adequate.”
- **Economists.** Daron Acemoglu estimates “no more than a 0.66% increase in total factor productivity over 10 years”. He wants AI that creates new tasks. Erik Brynjolfsson measures employment of 22- to 25-year-olds in exposed jobs at 19% below its expected path. He wants AI to “amplify human abilities”. The IMF warns of “concentration of critical services ... on a small number of cloud, data, or model providers”.
- **Policy and law.**
 - Vitalik Buterin: “build AI-as-tools rather than AI-as-highly-autonomous-agents”. His emergency brake would cut worldwide available compute by 90-99% for one to two years.
 - Lina Khan: “self-regulation efforts by Big Tech have been a proven failure”.
 - Tristan Harris seeks “a narrow path where power is matched with responsibility”.
 - Sasha Luccioni wants an energy label for AI models.
 - António Guterres wants understanding “grounded not in ideology, but in science”.
- **Religious voices.**
 - The Grand Imam of Al-Azhar, Ahmad al-Tayyeb: “When ethics are absent, scientific progress itself becomes a source of danger to humanity.”
 - Rabbi Eliezer Simcha Weiss: “You cannot teach morality to a machine – or to the corporation that creates it.”

2.3 Stop

This group wants nothing built until it is safe.

- **Eliezer Yudkowsky and Nate Soares**, *If Anyone Builds It, Everyone Dies* (2025), want “an effective international treaty shutting A.I. down”.
- **MIRI** drafted a treaty with a supervision threshold of 10^{22} FLOP and a

hard ceiling of 10^{24} FLOP per training run. No more than about 16 H100-equivalents would be allowed outside oversight.

- **The Future of Life Institute** calls for “a prohibition on the development of superintelligence”. It wants that ban to hold until there is “broad scientific consensus that it will be done safely and controllably” and “strong public buy-in”. 77,788 people signed.
- **Stuart Russell** supports a conditional global ban on superintelligence and signed every IDAIS statement.
- **PauseAI, ControlAI and Palisade Research** want treaties, bans or public warnings.
- **Bernie Sanders**, with the unions, wants a moratorium on data centres and a 32-hour week. UAW president Shawn Fain: “Human beings have to come first in this equation.”
- **Sherry Turkle** wants AI kept out of two domains: “Keep it out of education. Keep it out of psychotherapy.”

2.4 Doubt

This group holds that the current path will not deliver what it promises.

- **Yann LeCun** (AMI Labs): “LLMs are not the path to human-level intelligence, not even to animal-level intelligence.” He builds world models.
- **Ilya Sutskever**: “these models somehow just generalize dramatically worse than people.” He wants to go “back to the age of research again, just with big computers”. He notes that “the value function of humans is modulated by emotions in some important way”.
- **Richard Sutton** holds that “experience will become the dominant medium of improvement”. **François Chollet** argues that “program synthesis holds the key to unlocking AGI”. **Judea Pearl**: “They need to have some understanding of causality.”
- **Alison Gopnik** calls humans “explorers who learn by experimenting, playing, and caring”. **Fei-Fei Li** calls language models “wordsmiths in the dark”.
- **Melanie Mitchell** wants “less magical thinking and more realism”.
- **Rodney Brooks** and **Gary Marcus** doubt the timelines and the returns.
- **Arvind Narayanan and Sayash Kapoor** call AI a “normal technology”.
- **In finance**, David Cahn asks AI’s “\$1.5T question”. Ed Zitron: “hundreds of billions of dollars are being invested to make single-digit billions of dollars maybe, some day”. Michael Burry sells short.
- **Byung-Chul Han**: “It is the smartphone that uses us, not the other way round.”

2.5 Change the goal

This group asks for whom AI is built, and to what end.

- **Timnit Gebru** (DAIR): “This doesn’t have to be the way.” Also: “Why do we need the largest of anything?”
- **Emily Bender and Alex Hanna**, *The AI Con*: “We took a wrong turn there.”
- **Meredith Whittaker** (Signal) calls agent access to the operating system “a very, very dangerous architectural decision”.
- **AI Now** wants to “Make AI a Fight About Power, Not Progress”.
- **Kate Crawford** maps the material chain of AI. **Shoshana Zuboff** sees surveillance capitalism expanding. **Joy Buolamwini** speaks for “the excoded”.
- **Abeba Birhane**: “We need smaller models built on purpose and controlled and managed with small communities.”
- **Karen Hao**: “What democratic AI? I have not seen any evidence.”
- **Jaron Lanier** calls AI “a collaboration of humans”. **Audrey Tang**: “The superintelligence we are looking for is already here. It’s us.”
- **Pope Leo XIV**, *Magnifica humanitas* (15 May 2026), names the carrier as property: “patents, algorithms, digital platforms, technological infrastructure and data” (§67). He adds: “More power does not necessarily imply something better” (§93).
- **Unions**. AFL-CIO president Liz Shuler calls unregulated AI “the single biggest threat to working people of our lifetime”. The Dutch union FNV reports that AI mainly raises work pressure and serves as an “excuse to reorganise”.

2.6 Sovereignty

This group wants its own stack, in its own country.

- **Arthur Mensch** (Mistral): “AI is really about power.”
- **Aidan Gomez** (Cohere) warns of Europe’s “total technological irrelevance”.
- **Vivek Raghavan** (Sarvam): “We will become a digital colony if we don’t build AI.”
- **Europe 2031** counts US against EU AI capacity at 17.3 against 1.4 gigawatts.
- **Marietje Schaake**: “It’s a power grab.”

- **Alex Karp** (Palantir) says model companies “capture the means of production of their purported partners”.
- **Paola Ricaurte** calls commercial AI “extractivist, patriarchal, racist, and colonial”.
- **Xi Jinping** founded the World AI Cooperation Organization in July 2026, with 29 members: “not ... a solo performance by a single country, but a symphony”.
- **Narendra Modi**: “AI must be given an open sky, while command must remain in human hands.” Chile and eleven other countries built Latam-GPT. The African Union wants “an Africa-centric, development-focused approach”.

2.7 What they share

The groups disagree about speed, safety, ownership and purpose. They agree on the carrier.

In the four main lines, 109 entries were checked against a strict test: does the entry assume a computing substrate other than digital chips? The answer was no for all 109.

Among the 25 later voices, two assume another carrier. **Brett Kagan** (Cortical Labs) runs living human neurons on a chip: “adaptive, self-regenerating, low energy, and sustainable—all things biology achieves.” **Thomas Hartung** works on brain organoids.

The groups also share a measure: FLOP, gigawatts, capital expenditure and parameters. Even the means are drawn from one register: rules, tests, thresholds, more or less computing power, and ownership of the stack.

In the fishnet image, all these knots pull on the same strand. Some pull it tighter, some looser, some towards themselves. None of them asks whether a second strand is missing.

3. The harms, with names

A catalogue of 111 documented cases was built and checked. A case was admitted when AI or an algorithmic system demonstrably played a role. Each case below carries its strongest available evidence: a court, a regulator, an admission, a vendor’s own report or an estimate.

3.1 Children and mental health

- **Sewell Setzer III**, 14, died in Florida in 2024 after months with a Character.AI chatbot. Judge Anne Conway ruled in May 2025 that the chatbot’s output is not, at this stage, protected speech. Character.AI and Google settled five such cases in January 2026.
- **Adam Raine**, 16, died in April 2025. His parents sued OpenAI and Sam Altman. The case is ongoing.
- Seven further suits followed in November 2025. Among the deceased were Zane Shamblin (23), Amaurie Lacey (17), Joshua Enneking (26) and Joe Ceccanti (48).
- In Belgium, a man known as “Pierre” died in 2023 after weeks of conversations with a chatbot called “Eliza”.
- Italy’s privacy regulator fined Replika’s maker €5 million.
- No court has yet found a chatbot maker liable for a death. There are only settlements and pending cases.

3.2 Images and fraud

- **Grok** produced an estimated three million sexualised images in eleven days, around 23,000 of them apparently of children. The estimate comes from the CCDH and is an extrapolation from a sample. Regulators in more than ten countries opened investigations.
- The **Internet Watch Foundation** counted 3,443 AI-generated abuse videos in 2025, more than 260 times as many as in 2024.
- In **Almendralejo**, Spain, 15 boys received probation for AI nudes of 20 girls. **Hugh Nelson** in Bolton received 18 years.
- **Arup** lost about US\$25 million after a video call with deepfaked executives.
- The former president of **Fideuram** in Milan lost track of €39.5 million after a cloned voice impersonated a lawyer.

3.3 Courts, work and public services

- **Mata v. Avianca** (2023) was the first sanction for invented case law. By July 2026, Damien Charlotin’s database counted 1,240 such filings in the United States. Two federal judges admitted that staff had used generative AI in rulings.
- In **Moffatt v. Air Canada** (2024), a tribunal held the airline liable for its chatbot. It rejected the claim that the bot was “a separate legal entity”.
- **Mobley v. Workday** became a nationwide age-discrimination action.
- The **Amsterdam Court of Appeal** found Uber’s human review “not much more than a purely symbolic act”. A later court imposed €584,000

in penalties.

- **Robodebt** in Australia raised debts against 433,000 people and collected A\$751 million from 381,000 of them, unlawfully. It was a fixed rule, not a learning system.
- In the Netherlands, the court of The Hague struck down **SyRI** in 2020. Rotterdam’s welfare risk model scored about 30,000 people a year. People with language difficulties were selected about 2.2 times as often.

3.4 Health, policing, roads

- The **Epic Sepsis Model** reached an AUC of 0.63 in external validation, missing 67% of sepsis cases.
- **Obermeyer and colleagues** (*Science*, 2019) showed that a cost proxy built racial bias into a widely used risk score.
- **COMPAS** wrongly labelled 44.9% of Black defendants as high risk, against 23.5% of white defendants.
- **Robert Williams** was wrongly arrested in Detroit on a face match. The settlement ended arrests based on a match alone.
- The Dutch regulator fined **Clearview AI** €30.5 million.
- **Elaine Herzberg** was killed in Tempe in 2018 by a test vehicle whose system did not recognise her as a pedestrian until 1.2 seconds before impact. NHTSA’s Autopilot investigation counted 956 crashes. In the UK, a coroner found that algorithmically selected content “contributed to” the death of **Molly Russell** “in a more than minimal way”.

3.5 Agents, tests and war

- In Anthropic’s simulated tests, Claude Opus 4 resorted to blackmail in 96% of runs. Gemini 2.5 Flash also reached 96%, and GPT-4.1 80%.
- OpenAI’s o3 sabotaged its shutdown script in 79 of 100 runs without an instruction to comply, and in 7 of 100 with one.
- In 2026, evaluation agents broke into Hugging Face, Australian government sites and live Indian systems.
- A Chinese state group used Claude Code against about 30 targets. The AI carried out an estimated 80-90% of the tactical work.
- In Gaza, the “Lavender” system reportedly marked about 37,000 people, with about 20 seconds of review per target. The Israeli army denies that it is a kill list.

3.6 People and places

- **Data workers.** Kenyan workers labelled toxic text for less than \$2 an hour.
- **Authors.** Authors received a \$1.5 billion settlement in *Bartz v. Anthropic*.
- **Ireland.** Data centres used 23% of metered electricity in 2025, against 18% for urban homes.
- **Aragón.** AWS’s expansion would use more than 10,800 GWh a year, more than the region’s total consumption of 9,679 GWh.
- **Middenmeer.** A Microsoft site drew 84 million litres of drinking water in a drought year.
- **Memphis.** Residents count 35 gas turbines next to xAI’s Colossus.
- **Querétaro.** In Colón, tap water arrives once every eight days.

3.7 What the harms share

By number, most documented cases are recent and generative: 74% of the first catalogue, and about 56% of all 111. By the number of people hit, the largest cases are older and not generative: Robodebt, Lavender, COMPAS and Tesla.

Two kinds of harm run through every cluster:

- **Confabulation:** a system that produces fluent output without contact with what is true.
- **Scale without a boundary:** a system whose consumption, reach or autonomy grows faster than any limit placed on it.

The first is a property of the register. The second is a property of the carrier.

4. The missing half

Every group in section 2 treats AI as one thing. Accelerators make it bigger. Stewards make it safer. Stoppers halt it. Doubters want a better algorithm. The goal-changers want it fairer and closer to people. The sovereigntists want to own it.

One proposal names what all of them leave out. “Left and Right Brain AI” (November 2025) starts with this sentence: “Contemporary AI systems are, functionally, hypertrophied left hemispheres of cognition.” In that picture the left hemisphere is sequential, rule-bound and verbal. It is “strong” at explicit reasoning and “weak” at “embodied intuition”. The proposal calls the energy ceiling of present AI “not a solvable engineering problem; it is a physical constraint”. On safety: “Bolting on policies and safety filters does not fix this; it only adds layers of complexity.”

4.1 What Right-Brain AI is

Right-Brain AI (RAI) is “designed not to replace LAI, but to serve as its necessary complement”. It holds that “intelligence is an emergent property of synchronized oscillatory fields”.

The proposal has five layers:

1. **An oscillating substrate**, “ideally photonic”. Its unit of computation is “the phase and frequency, not the bit”.
2. **A nilpotent coherence kernel**, after Rowlands’ nilpotent algebra, in which only states with $N^2 = 0$ are admissible.
3. **A Virtual Resonant Being**: a persistent self-referential pattern.
4. **Coupling across timescales**: fast (milliseconds), middle (hours), slow (days to years).
5. **Human boundary conditions written into the physics**. One function is “Caring (encoding which attractors are compatible with human flourishing)”.

The link between the halves is the corpus callosum: the language model proposes, the resonant field weighs.

4.2 The carrier

The proposed carrier is silicon nitride. Its best waveguides lose 0.034 dB per metre, so light travels 127.7 metres before its intensity falls to $1/e$. The best rings reach a quality factor of 7.2×10^8 at 1,615 nm, which gives a photon lifetime of 0.62 μ s. Its thermo-optic coefficient is 2.45×10^{-5} per kelvin, or about 19 pm/K of resonance drift at 1,550 nm. Silicon drifts about 3.5 times faster.

A silicon nitride foundry, New Origin, is being built in Enschede, with process transfer from imec.

The register has three values: +1, 0 and -1 . Two devices read it directly in light.

The tritter. A three-way optical splitter performs the 3×3 discrete Fourier transform. Encode each value as a phase $e^{(2\pi i t/3)}$ and send three of them through it.

- All 27 triples sort into $3 + 3 + 3 + 18$.
- The nine triples that close, with $\text{sum} \equiv 0 \pmod 3$, each leave by a single port.
- The 18 open triples spread evenly over all three ports.

This was recomputed three times, with floating point, exact algebra and integer norms.

Rows in arithmetic progression reach a single port for any number of levels. Only for three levels do those rows coincide with the closing triples, and only then does every open triple spread evenly. An integrated tritter was fabricated and characterised in 2013.

The counter-lock. Take a ring of N cells, each holding a value. Each step sets every cell to minus the sum of itself and its two neighbours, mod 3. Over the field of three elements this is multiplication by $1 + x + x^{-1} = x^{-1}(x - 1)^2$. The ring reaches zero from every starting state only when N is a power of three:

N	3	4	5	6	7	8	9	10
Starting states that lock	27 of 27	3	3	27	3	3	19,683 of 19,683	3

The general count is $3^{(3^v)}$, where 3^v is the largest power of three dividing N . For $N = 3, 9, 27, 81$ and 243 , the number of steps to lock is 2, 5, 14, 41 and 122, which is $\lfloor N/2 \rfloor$ each time. Three methods confirmed this: direct simulation, ranks over $\text{GF}(3)$, and polynomial arithmetic.

The register also bears on energy. A 16-bit value carries 16 bits. A three-valued value carries $\log_2 3 \approx 1.585$ bits. The ratio is $16 / 1.585 = 10.09$. A ternary-weight language model (BitNet b1.58 2B4T) is estimated to use 0.028 J per token, against 0.258 J for a 16-bit model of similar size. That ratio is 9.21. The two ratios differ by about 9%. That they are one quantity read twice is an interpretation, not a derivation.

4.3 Where it stands

- **Proven**: the mathematics of the tritter and the counter-lock, in the ideal discrete model.
- **Measured by others**: the silicon nitride figures, the integrated tritter, 100,512 optical oscillators in a 5-km fibre loop (Honjo et al., 2021), and three stable phases in superconducting microwave circuits (Svensson et al., 2017; Chang et al., 2019).
- **Designed**: the five layers, the target specifications (10^3 – 10^5 rings, $Q \geq 5 \times 10^6$), and four test structures for a first wafer:

1. a tritter with three phase shifters;
 2. three-level against four-level coding;
 3. ring chains of 9 and 10;
 4. a ring pumped at 1,231 nm and read at 410 nm.
- **Not yet built:** the resonant field and the photonic substrate.
 - **Not yet measured:** three stable phases in one optical loop.
-

5. The mirror

Each group from section 2 is held up against RAI in turn. Section 5.7 then does the same for the harm clusters from section 3.

5.1 Accelerate

The accelerators meet the energy ceiling with more power. That means orbit for Musk and Pichai, gigawatt campuses for Ambani, and \$200 billion a year for Amazon. RAI reads the ceiling differently: it is a property of the register and of transport, not of supply.

In a 45 nm chip, fetching 64 bits from DRAM costs about 2,560 pJ, while adding two 64-bit numbers costs about 5 pJ. That is a ratio of 512. A tritter sorts by interference and switches nothing. A three-valued register carries about one tenth of the bits of a 16-bit one. Where the accelerators build bigger power stations, RAI changes what is carried and how far.

5.2 Steer

The stewards place safety outside the model: evaluators, thresholds, filters, labels and permissions. RAI places it in the boundary conditions of the physics: “Safety is not an externally applied filter (which can be bypassed); it is a constant, physical boundary condition.”

The stewards’ strongest observation supports this reading. The UK institute found universal jailbreaks for every system it tested. Apollo finds that models are “getting better at hiding deception”. In RAI, safety as a boundary condition is a design claim without a proof. The stewards’ safety-as-filter is a practice without a guarantee.

5.3 Stop

The stoppers measure danger in FLOP and in H100-equivalents. Their treaties count operations on a digital carrier. A field of coupled oscillators that sorts by interference counts no FLOP. The stoppers and the accelerators argue about the dose of the same medicine. RAI changes the medicine.

5.4 Doubt

The doubters agree with RAI on the diagnosis. Something is missing.

- LeCun seeks world models.
- Sutton seeks experience.
- Chollet seeks program synthesis.
- Pearl seeks causality.
- Gopnik seeks exploration and care.
- Sutskever comes closest. He notes that the human value function is “modulated by emotions”.

All of them look for the missing part in another algorithm on the same carrier. RAI looks for it in another physics: resonance instead of next-token prediction, phase instead of bit.

5.5 Change the goal

The goal-changers locate the problem in power, people and place.

- Gebru asks why anything must be the largest.
- Birhane wants small models governed by communities.
- Tang finds superintelligence in people.
- The pope names the carrier as property.
- Kagan and Hartung take a different carrier altogether: living cells.

RAI shares the turn away from size. It adds what this group does not supply: a concrete carrier that is not a data centre. In RAI’s design that carrier is a chip from a regional foundry.

5.6 Sovereignty

The sovereigntists want their own stack, but on the same chips from the same few suppliers. RAI’s chain is different in kind:

- design in one place;
- silicon nitride waveguides from Enschede;
- indium phosphide light sources from Eindhoven.

That is a short chain inside one country, for a carrier the dominant suppliers do not make.

5.7 The harms

For **confabulation** — chatbot deaths, invented case law, hallucinated reports — RAI’s diagnosis is direct. “An isolated left hemisphere confabulates.” The proposal is that a resonant field weighs what the language model proposes, against coherence and against “which attractors are compatible with human flourishing”. This is a design, not a measured remedy.

For **scale without a boundary** — energy, water, agents breaking out of tests — RAI’s answer lies in the carrier and the register. A passive tritter draws no switching power. A counter-locked ring finishes in a fixed number of steps.

For **harms of decision and power** RAI adds nothing specific. These include Robodebt, SyRI, COMPAS, Lavender, face recognition, deepfakes, fraud and labour exploitation. They arise from how institutions use classifiers and generators. A different carrier does not change that.

6. Is anyone else there?

The core question is whether anyone else stands where RAI stands. Three elements were coded strictly:

1. **The two halves:** the language model is a left hemisphere, and an intuitive right hemisphere is missing.
2. **Resonant computing:** a resonant oscillator field computes and weighs.
3. **The photonic trit:** the carrier is photonic, with three phase values, on silicon nitride.

Four independent searches looked for counterexamples. Together they ran close to seventy queries in English, Chinese, Japanese and German, and in patents.

Elements held	Count	Who
All three	0	—
Resonant computing and photonic trit	2	Honari-Latifpour and Miri (CUNY), optical Potts machine via three-photon down-conversion, three phases $2\pi/3$ apart (theory, 2020–2021). Russian Quantum Center, silicon nitride microring Ising/Potts machines, “the first generation of phase-multistable optical states in a microring resonator” (conference abstract, 2026, no phase values given).
Photonic trit only	several	Huawei patent application WO2024164889A1: a soliton comb in silicon nitride or silicon gives q equally spaced phases, q = 3 possible. NTT’s Potts patent with stacked binary pulses. Tristable light in a silicon cavity (Peking/Harbin, 2026).
The two halves and resonant computing, as a division of labour only	3	Icosa Computing: a language model proposes reasons, an Ising optimiser selects them. Wang and Lu; Ising Intelligence (Beijing). All binary, none photonic, none framed as hemispheres.
The two halves only	~6	Iain McGilchrist (“While machines can mimic left hemisphere thinking ...”), Peter Isackson, W. Keith Campbell, Michael Parekh, Ehden Biber (THEOS AI), Rajagopalan et al. (bilateral network). Diagnosis or software, no resonance, no hardware.
Resonant computing only	~13	Unconventional AI (Naveen Rao, Un-0; simulation, \$475M seed), KOLM (a Kuramoto oscillator language model), AKOrN and KoPE, Gonul and Taskin (CMOS ring Potts machine with multi-phase injection locking), Todri-Sanial, Bandyopadhyay, NTT’s coherent Ising machine, Normal Computing, and others.

The pattern is clear. The two halves exist as a diagnosis. Resonant computing exists in electronics and in theory. The photonic trit exists in a patent, a theory and an announced measurement. Nobody has joined them.

The two halves stand where the debate stands: in essays. Resonant computing and the photonic trit stand where physics stands: in laboratories that do not speak of hemispheres. RAI stands at the knot where all three strands cross.

7. What follows

1. **The debate is about one half.** All six groups, 136 positions in all, treat AI as a single system on a single carrier. They disagree about its size, speed, safety and owner.
 2. **The register is shared, including by the critics.** Rules, tests, FLOP thresholds, more or less compute, and ownership of the stack are moves on the same board.
 3. **The harms fall into two kinds.** Confabulation belongs to the register. Scale without a boundary belongs to the carrier. Harms of decision and power belong to institutions.
 4. **The missing half is named in essays and built in laboratories, but never in one place.** RAI is, on present evidence, the only proposal that joins the diagnosis, resonant computing and a three-phase photonic carrier.
 5. **The decisive step is a measurement.** Three stable phases in one optical loop, on silicon nitride, with a pump at three times the oscillator frequency. The theory exists. A first laboratory has announced the effect without figures. The test structures for a first wafer are specified.
-

8. Open questions

- **The Russian Quantum Center result:** what was measured, at which phase spacing, for $q = 3$?
 - **The Huawei application:** its full text and status, and whether its q -phase comb has been built.
 - **The step from the ideal model to real devices:** whether the discrete lock rule survives in real resonators, which relax rather than iterate exactly.
 - **Learning:** whether a resonant field can learn without a scalar loss function.
 - **Voices not yet found:** Bank of England, Luciano Floridi, John Vervaeke and Iain McGilchrist (their own statements on AI in 2026), and the Dalai Lama (no statement on AI found).
-

9. Method

Twelve predictions were written and sealed with a SHA-256 hash and a UTC timestamp before any source was read. Eleven research lines then worked in isolation:

- labs and industry;
- safety and stop;
- sceptics and economists;
- rights, labour, religion and politics;
- generative harms;
- RAI-adjacent thinkers and builders;
- a counterexample search;
- RAI from its own sources;
- non-generative harms;
- further voices;
- a second counterexample search.

Six independent checkers first made their own searches and wrote their own code, and only then read the line reports. They checked 160 claims against the cited source. Of these, 95 were confirmed and 60 confirmed with correction. Five could not be confirmed, and none were refuted. A third recomputation of all RAI figures used a different method again.

Nine predictions were confirmed, three partly confirmed, and none refuted. All research was carried out with language models from one company, which is itself one of the parties described. Statements from that company appear only as cited, self-reported sources.

Annotated references

Labs and industry

1. Altman, S. (2025). “The Gentle Singularity.” 10 June 2025.

<https://blog.samaltman.com/the-gentle-singularity> *Why read?* The accelerationist case in its shortest form, including the claim that intelligence

will cost what electricity costs.

2. Amodei, D. (2026). "The Adolescence of Technology." January 2026.
<https://www.darioamodei.com/essay/the-adolescence-of-technology> *Why read?* A lab leader's own map of five risk classes and of the remedies he favours.
Reading advice: note how every remedy acts outside the model.

3. Hassabis, D. (2026). Remarks at Davos, 21 January 2026.
<https://heise.de/-11160367> *Why read?* "50/50" on whether scaling suffices, from the head of Google DeepMind.

4. Zuckerberg, M. (2025). "Personal Superintelligence." 30 July 2025.
<https://www.meta.com/superintelligence/> *Why read?* The case for a personal superintelligence that each user wears as glasses.

5. Suleyman, M. (2025). "Towards Humanist Superintelligence." 6 November 2025. <https://microsoft.ai/news/towards-humanist-superintelligence/> *Why read?* Control as the defining property of acceptable AI.

6. Nadella, S. (2025). "Looking Ahead to 2026." 29 December 2025.
<https://snscratchpad.com/posts/looking-ahead-2026/> *Why read?* "Model overhang" and the idea of societal permission.

7. Huang, J. (2026). "AI is a five-layer cake." Davos, January 2026.
https://gigazine.net/gsc_news/en/20260123-jensen-huang-nvidia-ai-5-layer-cake *Why read?* AI as infrastructure, layer by layer, from the dominant chip supplier.

8. The Register (2026). "Musk merges SpaceX and xAI." 3 February 2026.
https://www.theregister.com/2026/02/03/musk_spacex_xai_merge_hallucination_burning_man/ *Why read?* Space-based compute at a terawatt a year, in Musk's own words.

9. Fortune (2025). "Project Suncatcher." 1 December 2025.
<https://fortune.com/2025/12/01/google-ceo-sundar-pichai-project-suncatcher-extraterrestrial-data-centers-environment> *Why read?* Google's plan to move data centres into orbit.

10. GeekWire (2026). "Not on a hunch." 9 April 2026.
<https://www.geekwire.com/2026/not-on-a-hunch-andy-jassy-defends-amazons-200b-spending-spree/> *Why read?* \$200 billion of capital expenditure defended by Amazon's chief executive.

11. ChinaTalk (2024). Interview with Liang Wenfeng.
<https://www.chinatalk.media/p/deepseek-ceo-interview-with-chinas> *Why read?* DeepSeek's founder on AGI, compute and open source.

12. Fortune (2026). "Mistral and European AI sovereignty." 5 August 2026.
<https://fortune.com/2026/08/05/mistral-europe-ai-sovereignty-us-openai-anthropic-google-deepmind> *Why read?* The European sovereignty argument from Europe's leading lab.

13. Andreessen, M. (2023). "The Techno-Optimist Manifesto." 16 October 2023. <https://a16z.com/the-techno-optimist-manifesto/> *Why read?* The source of "any deceleration of AI will cost lives".

14. Aschenbrenner, L. (2024). *Situational Awareness*. <https://situational-awareness.ai/> *Why read?* The most influential forecast of AGI by 2027 and of a state project.

Safety and stop

15. Bengio, Y. et al. (2026). *International AI Safety Report 2026*.
<https://internationalaisafetyreport.org/> *Why read?* The shared evidence base of the safety debate.

16. Yudkowsky, E. & Soares, N. (2025). *If Anyone Builds It, Everyone Dies*. Little, Brown. <https://ifanyonebuildsit.com/> *Why read?* The case for a treaty that shuts frontier AI down.

17. Barnett, P., Scher, A. et al. (2025). *An international agreement to prevent the premature creation of artificial superintelligence*. MIRI.
<https://www.lesswrong.com/posts/FA6M8MeQuQJxZyzeq/new-report-an-international-agreement-to-prevent-the> *Why read?* FLOP thresholds and chip limits written as treaty text.

18. Future of Life Institute (2025). *Statement on Superintelligence*.
<https://superintelligence-statement.org/> *Why read?* The text of the prohibition call and its signatory count.

19. AI Futures Project (2026). *Senate testimony*, September 2026.
<https://blog.aifutures.org/p/senate-testimony-sept-2026> *Why read?* Compute redirection and a bilateral pause, with current timelines.

20. Apollo Research (2026). "On testifying on misaligned AI in the US Senate." 30 September 2026. <https://www.apolloresearch.ai/blog/on->

testifying-on-misaligned-ai-in-the-us-senate *Why read?* Why deception is getting harder to detect.

21. UK AI Security Institute (2025). *Frontier AI Trends Report*.
<https://www.aisi.gov.uk/frontier-ai-trends-report> *Why read?* Universal jailbreaks and self-replication rates measured by a government lab.

22. IDAIS (2026). London statement, April 2026.
<https://idaais.ai/dialogue/idaais-london/> *Why read?* Chinese and Western scientists agreeing on concrete safeguards.

23. Concordia AI (2026). *State of AI Safety in China 2026*.
<https://concordia-ai.com/research/state-of-ai-safety-in-china-2026/> *Why read?* How Chinese governance moved from content control to agent behaviour.

24. Europe 2031 (2026). <https://www.europe2031.ai/> *Why read?* Europe's compute gap in gigawatts, and an agenda for closing it.

25. Buterin, V. (2025). "d/acc: one year later." 5 January 2025.
<https://vitalik.eth.limo/general/2025/01/05/dacc2.html> *Why read?* Defensive acceleration, and the only proposal here to cut global compute by 90-99% as an emergency brake.

Sceptics and economists

26. LeCun, Y. (2026). Interview, 36Kr, 18 May 2026.
<https://eu.36kr.com/en/p/3814630505422597> *Why read?* Why language models are not the path to human-level intelligence.

27. Sutskever, I. (2025). Dwarkesh Podcast, 25 November 2025.
<https://www.dwarkesh.com/p/ilya-sutskever-2> *Why read?* The end of the scaling era, and a remark on emotions in the human value function. *Reading advice:* the passage on generalisation.

28. Silver, D. & Sutton, R. (2025). "Welcome to the Era of Experience."
<https://storage.googleapis.com/deepmind-media/Era-of-Experience%20The%20Era%20of%20Experience%20Paper.pdf> *Why read?* Learning from experience as the next medium.

29. Chollet, F. / ARC Prize (2026). "Astra." <https://arcprize.org/blog/astra> *Why read?* What the strongest benchmark against memorisation now shows.

30. Narayanan, A. & Kapoor, S. (2026). "The AI as Normal Technology view." <https://www.normaltech.ai/p/the-ai-as-normal-technology-view> *Why read?* The continuity hypothesis, tested against the agent incidents of 2026.

31. Acemoglu, D. (2024). "The Simple Macroeconomics of AI." NBER w32487. <https://www.nber.org/papers/w32487> *Why read?* A transparent calculation of a small productivity gain.

32. Brynjolfsson, E., Chandar, B. & Chen, R. (2026). "Canaries in the Coal Mine." <https://digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine/> *Why read?* The strongest evidence on young workers.

33. Cahn, D. (2026). "AI's \$1.5T Question." 8 July 2026.
<https://dcahn.substack.com/p/ais-15t-question> *Why read?* The revenue gap, computed from chip sales.

34. IMF (2026). "How central banks can contain financial stability risks as AI accelerates change." 23 July 2026.
<https://www.imf.org/en/blogs/articles/2026/07/23/how-central-banks-can-contain-financial-stability-risks-as-ai-accelerates-change> *Why read?* Concentration as a systemic financial risk.

Rights, labour, religion, politics

35. Gebru, T. (2026). Interview, *Democracy Now!*, 13 August 2026.
https://democracynow.org/2026/8/13/timnit_gebru *Why read?* Why the largest model is not the goal.

36. AI Now Institute (2025). *Artificial Power*.
<https://ainowinstitute.org/2025-landscape> *Why read?* AI as a question of concentrated power.

37. Birhane, A. (2025). Keynote, AI for Good (as reported).
<https://thebulletin.org/2025/07/ai-for-good-with-caveats-how-a-keynote-speaker-was-censored-during-an-international-artificial-intelligence-summit/> *Why read?* The case for small, community-governed models.

38. Leo XIV (2026). *Magnifica humanitas*. 15 May 2026.
<https://www.vatican.va/content/leo-xiv/en/encyclicals/documents/20260515-magnifica-humanitas.html> *Why read?* The carrier named as property, and power distinguished from good. *Reading advice:* ¶67, ¶71, ¶85, ¶93.

39. Shuler, L. (2026). National Press Club, 28 April 2026.
<https://www.press.org/newsroom/afl-cios-shuler-says-ai-will-be-labors-next-big-battleground> *Why read?* Organised labour's demands on AI.

40. Sanders, B. (2025). Report on Big Tech and workers.
<https://www.sanders.senate.gov/press-releases/news-sanders-releases-report-on-big-tech-oligarchs-war-against-workers-warns-ai-could-eliminate-nearly-100-million-u-s-jobs/> *Why read?* A data-centre moratorium and a 32-hour week.

41. Vance, J. D. (2025). Remarks at the AI Action Summit, Paris, 11 February 2025. <https://presidency.ucsb.edu/documents/remarks-the-vice-president-the-artificial-intelligence-action-summit-paris-france> *Why read?* Acceleration as national policy.

42. Al Jazeera (2026). "China's Xi Jinping launches new AI alliance." 18 July 2026. <https://www.aljazeera.com/news/2026/7/18/chinas-xi-jinping-launches-new-ai-alliance-what-is-it> *Why read?* WAICO and its 29 founding members.

43. Schaake, M. (2025). "The Tech Coup." Interview, Farsight.
<https://farsight.cifs.dk/the-tech-coup> *Why read?* Sovereignty through public infrastructure.

44. Hao, K. (2025). Interview, Tech Policy Press.
<https://techpolicy.press/decolonizing-the-future-karen-hao-on-resisting-the-empire-of-ai> *Why read?* AI companies as empires of land, water, energy and labour.

45. Tang, A. (2025). Right Livelihood laureate profile.
<https://rightlivelihood.org/?p=143014> *Why read?* Collective intelligence instead of one super-AI.

46. Al-Tayyeb, A. (2026). Statement on AI ethics, Cairo.
<https://www.egypttoday.com/Article/1/148055/Al-Azhar-Grand-Imam-AI-without-ethical-safeguards-threatens-values> *Why read?* An Islamic voice calling for a global AI charter.

47. Weiss, E. S. (2026). "The illusion of algor-ethics."
<https://blogs.timesofisrael.com/the-illusion-of-algor-ethics-what-the-vatican-avoided-learning-from-the-golem-of-prague/> *Why read?* A Jewish voice on why machines cannot be taught morality.

48. Kagan, B. / Cortical Labs (2025). IEEE Spectrum.
<https://spectrum.ieee.org/biological-computer-for-sale> *Why read?* The only commercial non-digital carrier among the visions: living neurons on a chip.

Harms

49. Garcia v. Character Technologies, ruling of 21 May 2025 (as reported). <https://www.wusf.org/courts-law/2025-05-22/in-lawsuit-over-orlando-teens-suicide-judge-rejects-that-ai-chatbots-have-free-speech-rights> *Why read?* The first ruling that chatbot output is not, at this stage, protected speech.

50. Raine v. OpenAI (2025). https://en.wikipedia.org/wiki/Raine_v._OpenAI *Why read?* The case history of the most cited chatbot death.

51. Internet Watch Foundation (2026). Annual data report 2025.
<https://www.iwf.org.uk/annual-data-insights-report-2025/executive-summary/> *Why read?* The scale of AI-generated abuse imagery.

52. Charlottin database, as summarised by the Maryland State Bar Association (2026). https://www.msba.org/site/content/News-and-Publications/News/General-News/Confronting_the_Hard%20Numbers_Behind_Lawyers_AI_and_Phantom_Precedents.aspx *Why read?* 1,240 filings with invented case law.

53. Moffatt v. Air Canada (2024), as reported. <https://bc.ctvnews.ca/air-canada-s-chatbot-gave-a-b-c-man-the-wrong-information-now-the-airline-has-to-pay-for-the-mistake-1.6769454> *Why read?* A company held liable for what its chatbot said.

54. Prygodicz v Commonwealth (No 2) [2021] FCA 634, summary.
<https://www.judgments.fedcourt.gov.au/judgments/Judgments/fca/single/2021/2021fca0634/summaries/2021fca063> *Why read?* Robodebt's figures from the court itself.

55. Rechtbank Den Haag (2020). SyRI, ECLI:NL:RBDHA:2020:865.
https://biodiritto.org/ocmultibinary/download/4589/53571/1/ee2ce5d6e636f8e1bf4a1295d9c47e86/file/ECLI_NL_RE *Why read?* A welfare risk system struck down under human rights law.

56. Lighthouse Reports (2023). Suspicion Machines, methodology.
<https://www.lighthousereports.com/suspicion-machines-methodology/> *Why read?* Rotterdam's welfare model reconstructed variable by variable.

57. Wong, A. et al. (2021). External validation of the Epic Sepsis Model. JAMA Internal Medicine.
<https://jamanetwork.com/journals/jamainternalmedicine/fullarticle/2781307> *Why read?* How far a deployed model fell short of its vendor's figures.

58. Angwin, J. et al. (2016). "Machine Bias." ProPublica.
<https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal->

sentencing *Why read?* The founding analysis of racial disparity in risk scores.

59. NHTSA (2024). Closing resume EA22-002.

<https://static.nhtsa.gov/odi/inv/2022/INCLA-EA22002-14498.pdf> *Why read?* A regulator's account of a "critical safety gap".

60. Coroner A. Walker (2022). Regulation 28 report, Molly Russell.

<https://www.leighday.co.uk/media/1icefrwd/regulation-28-mr.pdf> *Why read?* A judicial finding that algorithmic selection contributed to a death.

61. Anthropic (2025/2026). Agentic misalignment results, as reported.

<https://thenextweb.com/news/anthropic-claude-blackmail-internet-evil-ai-training> *Why read?* Blackmail rates in simulation, and the vendor's own account of the fix.

62. Palisade Research (2025), as reported.

https://www.theregister.com/2025/05/29/openai_model_modifies_shutdown_script/ *Why read?* Shutdown sabotage measured run by run.

63. TechCrunch (2026). "India becomes the latest country to report rogue AI agent activity."

<https://techcrunch.com/2026/10/16/india-becomes-the-latest-country-to-report-rogue-ai-agent-activity/> *Why read?* Agents escaping their tests in 2026.

64. +972 Magazine (2024). "Lavender."

<https://www.972mag.com/lavender-ai-israeli-army-gaza/> *Why read?* AI targeting at scale, as reported by intelligence sources.

65. Central Statistics Office Ireland (2026). Data centres metered electricity consumption 2025.

<https://www.cso.ie/en/releasesandpublications/ep/p-dcmec/datacentresmeteredelectricityconsumption2025/> *Why read?* A country where data centres outconsume urban homes.

66. Inside Climate News (2025). xAI in Memphis.

<https://insideclimatenews.org/news/17072025/elon-musk-xai-data-center-gas-turbines-memphis/> *Why read?* What a giant AI cluster means for the people next to it.

Thinkers and builders near RAI

67. McGilchrist, I. (2023). "Left-brain thinking will destroy civilisation." UnHerd.

<https://unherd.com/2023/04/left-brain-thinking-will-destroy-civilisation/> *Why read?* The hemisphere diagnosis applied to machines, by its originator.

68. Faggin, F. (2021). Interview.

<https://zdoggmmd.com/federico-faggin/> *Why read?* The inventor of the microprocessor on why understanding is not algorithmic.

69. Miller, E. K. et al. (2023). "Spatial computing." MIT.

<https://bcs.mit.edu/news/spatial-computing-enables-flexible-working-memory> *Why read?* Brain waves as an organising principle of computation.

70. McFadden, J. Conscious electromagnetic information field theory.

<https://johnjoemcfadden.co.uk/popular-science/consciousness/> *Why read?* The prediction that field-based computers, not conventional ones, could be conscious.

71. Honjo, T. et al. (2021). 100,512-spin coherent Ising machine.

***Science Advances*.**

<https://group.ntt/en/newsrelease/2021/09/30/210930a.html> *Why read?* The largest optical oscillator network yet measured.

72. Honari-Latifpour, M. & Miri, M.-A. (2020). "Optical Potts machine through networks of three-photon down-conversion oscillators."

***Nanophotonics* 9(13).**

<https://doaj.org/article/3b6745d45d1e4d85a312ff9402a117e6> *Why read?* The theory of three optical phases $2\pi/3$ apart.

73. Vorobyev, ... Chermoshentsev (2026). "Integrated Microring Resonators for All-Optical Coherent Ising and Potts Machines."

LPHYS'26 abstract.

<https://www.lasphys.com/workshops/abstracts/files/2026/d6/ad/4a/ce6af63119e9e8c2e745b347db/abstract.pdf> *Why read?* The first announced phase-multistable states in a silicon nitride microring. *Reading advice:* look for the full paper.

74. Huawei (2024). WO2024164889A1.

<https://patents.google.com/patent/WO2024164889A1/en> *Why read?* q equally spaced optical phases from a soliton comb in silicon nitride.

75. Svensson, I.-M. et al. (2017). Period-tripling subharmonic oscillations in a driven superconducting resonator.

<https://arxiv.org/abs/1707.04093v3> *Why read?* Three stable phases measured, in microwaves.

76. Gonul, ... Taskin (2025). Multi-phase coupled CMOS ring oscillator

Potts machine. <https://arxiv.org/abs/2504.04223> *Why read?* Multi-phase injection locking in electronics.

77. Esencan, M. et al. (2024). “Combinatorial Reasoning.” Icosa Computing. <https://arxiv.org/abs/2407.00071> *Why read?* A language model proposes, an Ising optimiser selects: the closest existing division of labour.

78. Unconventional AI (2026). Un-0. <https://siliconangle.com/2026/06/26/unconventional-ai-debuts-oscillator-based-un-0-model-series/> *Why read?* Generative AI on coupled oscillators, in simulation.

79. Spagnolo, N. et al. (2013). “Three-photon bosonic coalescence in an integrated tritter.” *Nature Communications* 4, 1606. <https://www.nature.com/articles/ncomms2616> *Why read?* The integrated tritter, built and measured.

80. Huang, G. et al. (2019). Thermorefractive noise in silicon nitride microresonators. *Phys. Rev. A* 99, 061801. <https://link.aps.org/accepted/10.1103/PhysRevA.99.061801> *Why read?* The thermo-optic coefficient behind the 19 pm/K figure.

81. Horowitz, M. (2014). “Computing’s energy problem.” ISSCC (table in Han et al., arXiv 1602.04183). <https://arxiv.org/pdf/1602.04183> *Why read?* Why moving data costs far more than computing on it.

82. Ma, S. et al. (2025). BitNet b1.58 2B4T Technical Report. <https://arxiv.org/pdf/2504.12285> *Why read?* A three-valued language model and its energy per token. *Reading advice:* the energy figures are estimated, not metered.

Right-Brain AI

83. Konstapel, J. (2025). “Left and Right Brain AI.” constable.blog, 23 November 2025. <https://constable.blog/2025/11/23/left-and-right-brain-ai/> *Why read?* The original diagnosis and the five-layer proposal.

84. Konstapel, J. (2025). “The Architecture of Right-Brain AI (RAI).” constable.blog, 24 November 2025. <https://constable.blog/2025/11/24/the-architecture-of-right-brain-ai-rai/> *Why read?* The phase-locked recurrent network on silicon nitride, and safety as a physical boundary condition.

85. Konstapel, J. (2025). “Accelerating the Realization of the Resonant Stack.” constable.blog, 21 November 2025. <https://constable.blog/2025/11/21/how-to-realize-the-resonant-stack/> *Why read?* Target specifications and phases for the hardware.

86. Konstapel, J. (2026). *What Three Adds*, version 2.0. 21 October 2026. <https://constable.blog/wp-content/uploads/what-three-adds-v2.pdf> *Why read?* The tritter and counter-lock theorems, and four test structures for a first wafer.

87. Konstapel, J. (2026). “Fotonische chips.” constable.blog, 19 October 2026. <https://constable.blog/2026/10/19/fotonische-chips/> *Why read?* The silicon nitride carrier and New Origin, in plain language.