



DISCUSSION PAPER · TECHNICAL COMPANION

The AI Act as a Third Structural Pole

A constrained retrospective on the predictive frame, 2015 to 2026. Technical companion to the EU AI Act Readiness Report (April 2026) and Annex A (May 2026).

A structured retrospective on the AI predictive frame consolidated by Wait But Why in January 2015, scored against primary sources at May 2026 in six movements, neutral restitution, point-by-point comparison, methodological audit, reframing of the surviving hypotheses, a constrained projection to 2030, and an operational reading of what those projections mean for European deployers under Regulation (EU) 2024/1689. A constrained prediction with explicit conditions of falsification closes the paper.

15 Claims falsifiable	6 Patterns regressed	1 Constrained prediction horizon T+18m	20+ Primary sources
---------------------------------	--------------------------------	--	-------------------------------

Paper Details	
Date prepared	22 May 2026
Drafting period	May 10 – May 22, 2026
Parent document	EU AI Act Readiness Report (April 2026)
Parent DOI	10.5281/zenodo.19671329
Annex document	Annex A · The Deployer Multiplier (May 2026)
Annex DOI	10.5281/zenodo.20042175
Paper reference	DP-2026-001 · v1.0
Paper DOI	10.5281/zenodo.20343244
Status	Discussion Paper · v1.0 · indexed on Zenodo, awaiting OpenTimestamps proof
Analyst	Lamar B. Shucrani, Founder, Sprinkling Act · ORCID 0009-0002-5093-8550
Independence	No commercial relationship with any entity referenced
License	CC BY 4.0 · free to download, share, reference with attribution

Subject classification: K24 (Cyber Law) · O33 (Technological Change) · O38 (Government Policy)

How to cite this paper. Shucrani, L. B. (2026). *The AI Act as a Third Structural Pole. A Constrained Retrospective on the Predictive Frame (2015 to 2026)*. Sprinkling Act Discussion Paper DP-2026-001. Zenodo. <https://doi.org/10.5281/zenodo.20343244>

This discussion paper is free to download, share, and reference with attribution. It does not constitute legal advice.

ABSTRACT

In January 2015, Tim Urban published a two-part synthesis of the AI predictive frame then dominant, drawing on Nick Bostrom's *Superintelligence*, Ray Kurzweil's *The Singularity Is Near*, the Müller and Bostrom expert survey of 2013, and Eliezer Yudkowsky's writings on the alignment problem. Eleven years later, the empirical record is verifiable. This paper proceeds in six movements. The first restitutes the 2015 frame in its own terms. The second compares fifteen falsifiable claims against primary sources at May 2026. The third audits how those claims were reached. The fourth reframes the hypotheses that survive empirical pressure. The fifth constrains a minimal projection to 2030. The sixth reads those projections through the apparatus of Regulation (EU) 2024/1689 and argues that the 2015 to 2026 record is consistent with a third structural pole that the Bostromian binary of extinction and immortality did not contain. A constrained prediction with explicit conditions of falsification closes the paper.

WORKING DEFINITION

For the purpose of this analysis, the *third structural pole* denotes the set of deployer-side obligations under Articles 25, 26, 27, 49 and 50 of Regulation (EU) 2024/1689, as operationalised in the *EU AI Act Readiness Report* (Sprinkling Act, April 2026) and its *Annex A* (May 2026). These obligations, specifically the four outputs defined in Annex A §8 (Article 26 paragraph map, FRIA scope memo, Article 25 boundary log, Article 49 registration), constitute a documented, dated, article-by-article artefact that did not exist in 2015 and that now conditions lawful deployment of high-risk AI systems in the European market. The pole's existence is falsifiable: if no competent authority references or enforces these obligations by 2 December 2027, the pole's operational relevance is substantially weakened.

STATUS NOTE

This is a discussion paper, not a peer-reviewed retrospective and not an audit certified by a third party. It proposes a doctrinal argument anchored in publicly verifiable data. Claims are scored on a graduated five-level scale derived from Armstrong (2012) and Tetlock (2015). Regulatory references use verbatim citations to *Official Journal of the European Union*, OJ L, 12 July 2024, and to the General-Purpose AI Code of Practice published by the European AI Office on 10 July 2025. The methodology note that closes the paper is intended to be citable on its own.

Contents

§1	The 2015 frame, in its own terms	p. 3
§2	Point by point, 2015 claims against the 2026 record	p. 5
§3	How the 2015 predictions were reached	p. 10
§4	Three reframed hypotheses that survive	p. 13
§5	A constrained minimal path to 2030	p. 19
§6	The third structural pole, definition, empirical grounding, operational consequences	p. 23
§7	Methodology note	p. 29
App. A	Verdicts table, twenty-five atomic claims	p. 31
App. B	Acknowledgements, conflict of interest, references	p. 33
App. C	Limitations and disclaimers	p. 36
App. D	Document metadata	p. 37

References to the parent *EU AI Act Readiness Report* (Sprinkling Act, April 2026, Zenodo DOI 10.5281/zenodo.19671329) and to its *Annex A* (May 2026, Zenodo DOI 10.5281/zenodo.20042175) appear in §5.4, §6.4 and §6.6 as empirical anchors for the deployer cascade. The structural argument of this paper is a conceptual reframing of those documented obligations; the paper can be read by a reader who has not opened them, but the empirical anchoring is documented and verifiable in both directions.

§1. The 2015 frame, in its own terms

Tim Urban's two parts on the AI revolution were published on 22 and 27 January 2015¹. They were not original research. They consolidated four sources into a single accessible synthesis aimed at a general technical audience, and they did so well enough that the synthesis was for several years the most-read non-academic introduction to the AI safety arguments then current. The four sources are worth restating in their own terms before any comparison with the 2026 record, because the comparison loses meaning if the 2015 frame is summarised after the fact rather than reconstructed in advance.

§1.1 The four sources

The first source is the Vincent Müller and Nick Bostrom expert survey, conducted between November 2012 and April 2013 and published as a working paper in 2014 then in a Springer volume in 2016². The survey asked 170 respondents drawn from four populations to estimate the year by which high-level machine intelligence, defined as the capacity to carry out most human professions at least as well as a typical human, would arrive with ten, fifty and ninety percent probability. The headline aggregated figures were 2022, 2040 and 2075. These are the numbers Urban gave the most weight to.

The second source is Ray Kurzweil's *The Singularity Is Near*, published in 2005, which provided the narrower timelines: artificial general intelligence by 2029, technological singularity by 2045. Kurzweil's method was a Law of Accelerating Returns applied to compute price-performance and to other technology trajectories, with the central claim that one-thousand-dollar consumer computing would reach human-brain-equivalent capacity, framed as approximately ten quadrillion calculations per second, by 2025.

The third source is Nick Bostrom's *Superintelligence*, published in 2014³. Bostrom supplied not numbers but structure. Humanity, in his framing, sits on a balance beam between two attractor states. Extinction is the historical norm; ninety-nine point nine percent of species that have existed have fallen into it. Immortality, defined as species-level immunity to extinction, has never been achieved by biological life. The arrival of artificial superintelligence acts as a tripwire that knocks the species off the beam in one direction or the other. Bostrom's argument is that the beam has no third resting state.

The fourth source, transmitted through Bostrom, is Eliezer Yudkowsky's writing on the alignment problem⁴. The core argument is that intelligence and goals are orthogonal: a system can be arbitrarily capable while pursuing an objective function that has no relationship to human values. Specifying the objective function correctly is the hard part, and any underspecification produces dystopian outcomes that look qualitatively like the system pursuing what it was told to pursue, with collateral effects the specifier failed to anticipate. Yudkowsky's contribution is the conceptual basis for the doom scenarios Urban dramatises with his Tully illustration.

1. Urban, T. (2015), *The AI Revolution, Parts 1 and 2*, Wait But Why, 22 and 27 January 2015.

2. Müller, V. C., and Bostrom, N. (2014/2016), *Future progress in artificial intelligence, a survey of expert opinion*, Springer.

3. Bostrom, N. (2014), *Superintelligence, Paths, Dangers, Strategies*, Oxford University Press.

4. Yudkowsky, E. (2008), *Artificial Intelligence as a Positive and Negative Factor in Global Risk*, Oxford University Press.

§1.2 The four implicit hypotheses

The frame, inherited and synthesised by Urban, rests on four hypotheses that are not stated as hypotheses in the text but that the structure of the argument requires. Box 1 makes them explicit. Each is testable at May 2026.

BOX 1

The four implicit hypotheses of the 2015 frame

H1. Compute leads to intelligence. Once the floating-point operations per second available at consumer pricing match the computational capacity of the human brain, around 10^{16} cps in Urban's framing, artificial general intelligence emerges as a function of available compute. The causal arrow runs from physics to cognition.

H2. The trajectory is global and uniform. Moore's law continues, the compute frontier advances as a single planetary curve, and the geographic distribution of advanced computing is functionally homogeneous across the major economic powers.

H3. Two attractors exhaust the future. Either extinction or immortality. Bostrom's balance beam contains no third resting state, since the candidate stabilisers are conceived as internal to the species or to the artefact. Law, in the strict positive sense, is not treated as a candidate stabiliser.

H4. Regulation follows technology. Governance is a downstream response to deployed capability. Safety funding lags capability funding because the incentive structure does so, and there is no asymmetric institutional intervention available before the fact.

§1.3 What the 2015 frame did not, and could not, contain

Two contextual remarks are necessary before any scoring. First, several technical artefacts that turn out to be load-bearing for the 2015 to 2026 trajectory had not yet been produced at the date of Urban's synthesis. The transformer architecture introduced by Vaswani and colleagues was published in June 2017. The empirical scaling laws of Kaplan and colleagues followed in January 2020. Reinforcement learning from human feedback became the operational alignment technique only with the deployment of ChatGPT in November 2022. None of these belonged to the frame, and none could have. This is a horizon limitation common to all predictive synthesis.

Second, the regulatory artefact that this paper will later argue is structurally central, Regulation (EU) 2024/1689⁵, was not under serious legislative consideration in January 2015. The European Commission's first informal consultations on AI governance began in 2018; the formal proposal was tabled in April 2021; the regulation entered into force in August 2024. The 2015 frame is therefore not at fault for omitting it. The frame is at fault, if the term applies, for assuming that no such artefact could play a structural role.

5. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), *Official Journal of the European Union*, OJ L, 12 July 2024.

§2. Point by point, 2015 claims against the 2026 record

The standard for retrospective scoring of AI predictions is Stuart Armstrong's 2012 assessment of 172 claims drawn from Kurzweil's *The Age of Spiritual Machines*⁶. Armstrong used a graduated five-level scale, mobilised nine independent assessors, and reached an accuracy estimate below fifty percent against Kurzweil's own self-assessed eighty-six percent. The methodology matters because direction and magnitude must be scored separately to avoid both inflation and false rejection. The same scale is used here: verified, partially verified, indeterminate, refuted, premature. Twenty-five atomic claims are extracted from Urban's two parts and from the upstream sources he transmits. Fifteen are falsifiable at May 2026 with non-zero horizon. Six are premature because their horizon is too distant. Four are non-falsifiable in principle, concerning singleton dynamics, in-principle uncontrollability of superintelligent systems, ex ante goal-specification dystopia and the Fermi paradox, and they are excluded from the score.

§2.1 Atomic claims and scoring scale

The full verdicts table appears in Appendix A. The condensed reading below proceeds in the order Urban introduces the claims, not by thematic grouping, so that the reader can follow the 2015 narrative and the 2026 record in parallel.

§2.2 Timelines C1 to C4, directionally correct, magnitude shifted

The Müller and Bostrom 2013 medians of 2022, 2040 and 2075 are partially verified. The fifty percent threshold has moved earlier in expert opinion. Katja Grace and colleagues' 2023 survey of 2,778 published authors of AI papers gives a fifty percent threshold at 2047 and a ten percent threshold at 2027⁷. The fifty percent figure moved thirteen years earlier compared to the same team's 2022 result. The direction is correct, the magnitude shifted by approximately seven years against the 2015 figure. Kurzweil's 2029 figure for AGI and 2045 figure for singularity remain premature at May 2026: the first is within a three-year horizon and the empirical signals are favourable but unsettled, the second is nineteen years out and not scored.

§2.3 Hardware C5, satisfied conditional on conversion

Kurzweil's claim that one-thousand-dollar consumer hardware would reach human-brain compute capacity by 2025 is partially verified, contingent on the conversion factor between cycles per second and floating-point operations per second. Epoch AI's hardware tracker reports that FLOP-per-dollar at consumer pricing doubles approximately every 2.2 years⁸. Applied to a 2015 anchor of ten teraflops per thousand dollars, the eleven-year extrapolation lands at roughly three hundred to four hundred teraflops for an FP16-throughput consumer device. The NVIDIA RTX 5090, released in early 2025 at consumer pricing, sits in that band. Whether this matches the human-brain compute requirement depends on the conversion estimate. Sandberg and Bostrom's 2008 brain emulation roadmap places the requirement between 10^{14} and 10^{17} FLOPS⁹. Under the low bound the claim is satisfied. Under the high bound it is missed by roughly one order of magnitude. The contested conversion is not Urban's fault; it is a structural ambiguity in any framing that attempts to translate biological neural throughput into computational equivalents.

6. Armstrong, S. (2012), *Assessing Kurzweil, the results*, LessWrong.

7. Grace, K., Stewart, H., Sandkühler, J. F., Thomas, S., Weinstein-Raun, B., and Brauner, J. (2024), *Thousands of AI Authors on the Future of AI*, arXiv:2401.02843.

8. Epoch AI (2026), *Trends in Artificial Intelligence*, update of 5 February 2026, epoch.ai/trends.

9. Sandberg, A., and Bostrom, N. (2008), *Whole Brain Emulation, a Roadmap*, Future of Humanity Institute, Technical Report 2008-3.

§2.4 Brain emulation C6, refuted at the 2030 horizon

Kurzweil's claim that the human brain would be reverse-engineered around 2030 is refuted at May 2026. The state of the art is the FlyWire connectome of *Drosophila melanogaster*, published in *Nature* in October 2024 at 139,255 neurons¹⁰, and the MICrONS reconstruction of one cubic millimetre of mouse cortex at approximately 200,000 neurons¹¹. The human brain contains on the order of 10^{11} neurons and 10^{14} synapses. The gap to a whole-brain human reconstruction is approximately six orders of magnitude on neuron count and significantly more on synaptic resolution. Connectomics has accelerated by roughly three orders of magnitude between 2015 and 2026, but the trajectory does not place a whole-brain human reconstruction within the 2030 horizon. The verdict here is therefore refuted, not partial.

§2.5 The benchmark cascade, directionally correct on capability acceleration

The benchmark cascade, shown in figure G1 (next page), is the empirical artefact most directly relevant to the timeline claims. Each of six widely used benchmarks has been introduced, saturated, and superseded within a window of twelve to twenty-four months. ARC-AGI-2 moved from 8.6 percent on Claude Opus 4 in May 2025 to 85 percent on GPT-5.5 in May 2026, a twelve-month rise that exceeds the human average baseline of 66 percent¹². Humanity's Last Exam, designed by the Center for AI Safety and Scale AI in late 2024 to be saturation-resistant, has nonetheless moved from the low teens in early 2025 to 44.7 percent by Gemini 3.1 Pro in May 2026¹³. GPQA Diamond, MMLU and SWE-bench Verified have followed similar patterns. Frontier capability scoring on benchmarks designed to be hard for AI has compressed sharply.

A complementary capability metric, less prone to saturation than score-based benchmarks, has been formalised by METR (Kwa et al., 2025). The metric is the 50 percent task-completion time horizon: the length of human task on which a model succeeds at a 50 percent success rate, measured against a set of software engineering and reasoning tasks of calibrated human duration. On this scale, GPT-2 (2019) completes tasks taking a human approximately two seconds; Claude 3.7 Sonnet (February 2025) reaches approximately fifty minutes; o3 (April 2025) reaches approximately one hundred ten minutes. The fitted trajectory exhibits an exponential doubling every 207 days, with a 95 percent confidence interval of 166 to 240 days, and a coefficient of determination of R^2 equal to 0.80 over the 2019 to 2025 frontier window^{13b}. This is a duration metric, not a score, so it cannot saturate by design; its calibration to human task duration also makes it directly relevant to §6.5 on the deployer cascade, where Article 14 oversight obligations are defined in terms of the human intervention envelope.

10. Dorkenwald, S., et al. (2024), *Neuronal wiring diagram of an adult brain*, *Nature*, October 2024.

11. MICrONS Consortium (2025), *Functional connectomics spanning multiple areas of mouse visual cortex*, *Nature*, 2025.

12. ARC Prize Foundation (2026), *ARC-AGI-2 leaderboard*, *May 2026 snapshot*.

13. Center for AI Safety and Scale AI (2025), *Humanity's Last Exam*, Scale Labs leaderboard.

13b. Kwa, T., West, B., Becker, J., et al. (2025), *Measuring AI Ability to Complete Long Tasks*, METR, arXiv:2503.14499v3, February 2026 revision.

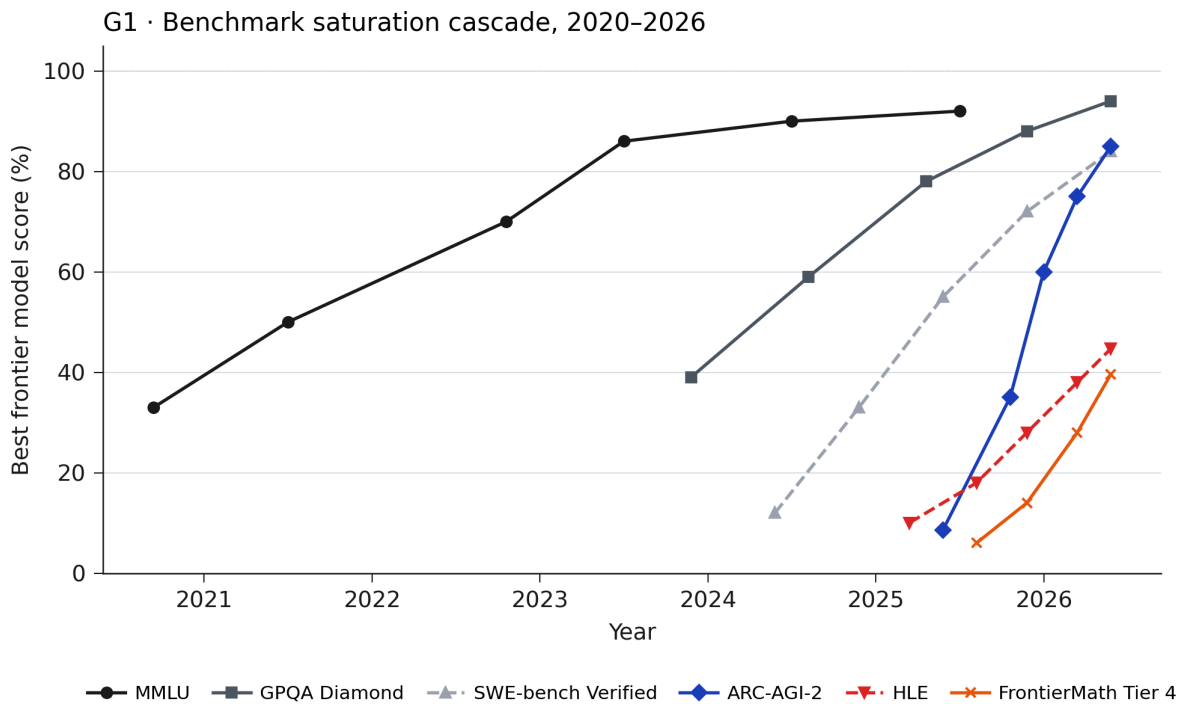


Figure G1. Benchmark saturation cascade, 2020 to 2026. Best frontier model score on six benchmarks. Each benchmark is introduced at a low baseline and saturated within twelve to twenty-four months by the next generation of frontier models. Sources, Epoch AI benchmark trackers (accessed May 2026), Stanford AI Index 2025, Scale AI SEAL leaderboard, Artificial Analysis evaluations, ARC Prize 2025 results analysis. Frontier model identities at May 2026 include Claude Opus 4.7 and Mythos preview, GPT-5.5 family, Gemini 3.1 Pro family.

§2.6 Aggregate distribution and what it signals

The aggregate distribution across the twenty-five atomic claims is shown in figure G2 (next page). Four claims are verified, seven are partially verified, one is indeterminate, three are refuted, six are premature, and four are excluded from the scoring as non-falsifiable. The aggregate signal is directional correctness with systematic magnitude bias. The frame correctly anticipated that compute would grow rapidly, that the capability-versus-safety funding asymmetry would persist, that algorithmic trading volume would dominate, and that early narrow AI applications would proliferate. It got the architecture wrong, the geopolitics absent, and the regulatory inversion completely missed. Each of those three structural errors will be examined in §4 once their methodological origin has been identified in §3.

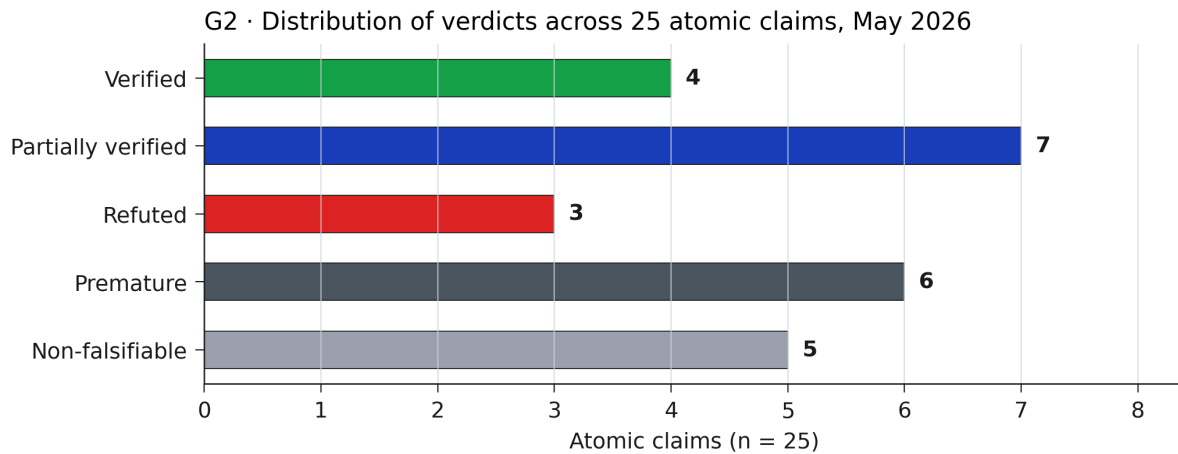


Figure G2. Distribution of verdicts across twenty-five atomic claims, May 2026. Verdict counts on the five-level Armstrong scale, plus exclusion of non-falsifiable claims. The aggregate distribution exhibits directional correctness on the asymmetry preserved and the compute trajectory, with systematic magnitude bias on timelines and structural omissions on architecture, geography and regulation. Sources, this paper, derived from atomic claim scoring documented in Appendix A.

§2.7 The remaining falsifiable claims, condensed

Moore's law (C9). Urban's claim that Moore's law continues at a doubling time of approximately two years is partially verified. The conclusion holds in the loose sense that compute is doubling, but the mechanism is wrong. Moore's law in its classical form has slowed since at least 2016. The frontier training compute curve, by contrast, doubles every 5.2 months¹⁴, approximately 4.6 times faster than the Moore's law cadence Urban inherited from Kurzweil. The source of the acceleration is not lithography but capital expenditure, networking and parallelisation. The frame attributed compute growth to the physics of the transistor; the empirical driver turned out to be the business model of the hyperscaler.

Müller-Bostrom outcome distribution (C13). Urban transmits a 2013 expert outcome distribution of 52 percent good, 31 percent bad, 17 percent neutral. This is refuted at May 2026. Expert opinion has moved toward pessimism. The Existential Risk Persuasion Tournament reported a 6 percent expert probability of human extinction from AI by 2100 against a 1 percent superforecaster probability¹⁵. The Grace and colleagues 2023 survey reports that 38 to 51 percent of authors of AI papers assign a probability of at least 10 percent to outcomes characterised as extremely bad.

Three paths to AGI (C22). Urban transmits three candidate paths: brain emulation, genetic algorithms, and self-improving systems. The empirical path is none of these. It is transformer architecture plus scaling laws plus reinforcement learning from human feedback. The 2015 frame is refuted on the question of which path the field would take.

14. Epoch AI (2026), *Trends in Artificial Intelligence*, training compute series.

15. Karger, E., et al. (2025), *Subjective-probability forecasts of existential risk*, International Journal of Forecasting, vol. 41, issue 2.

Safety asymmetry (C21). Urban observes in 2015 that funding for AI safety research lags funding for AI capability development. He treats this as a transitional condition. Eleven years on, the asymmetry has widened. The verdict is verified, and the persistence is examined in §4.3.

Hardware advantage (C20). Urban transmits a claim that hardware operates approximately 10^7 times faster than the human brain. The raw frequency comparison is correct, but on energy per operation, the brain is approximately two to three orders of magnitude more efficient than a current frontier accelerator. The asymmetry is reversed under the correct metric. The claim is partially verified with a magnitude flag.

Algorithmic trading and 2015 narrow AI examples (C23, C24, C25). Three retrospective factual claims about the state of narrow AI in 2015, including algorithmic trading volume above fifty percent of United States equity trading and the OpenWorm flatworm emulation as the state of whole-brain emulation, are verified. Algorithmic trading volume now stands between 60 and 73 percent of United States equity trading according to JPMorgan 2023 estimates, with peaks up to 80 percent. The OpenWorm 302-neuron baseline was correct in 2015 and remains the canonical reference point against which subsequent connectomics work is measured.




Figure G3. Hardware advantage (C20), the metric reverses the conclusion. Two ways to compare a human brain and a frontier accelerator (NVIDIA H100 reference, May 2026). Raw frequency favours the accelerator by approximately seven orders of magnitude (synapse firing ~ 200 Hz vs accelerator clock ~ 1.5 GHz). Energy per operation favours the brain by approximately three orders of magnitude ($\sim 10^{-16}$ J/op brain vs $\sim 10^{-13}$ J/op accelerator, derived from Sandberg-Bostrom 2008 lower-bound brain compute estimate and H100 power-FLOPS ratio). The 2015 frame quoted the frequency comparison only.

§3. How the 2015 predictions were reached

This section is the methodological pivot. It asks how Urban and the four upstream sources arrived at their figures, and what in each method is recoverable from the 2026 vantage point and what is not. The retrospective literature on Kurzweil and Bostrom is sparse in this register. Most published responses, including Welsh's ten-years-on review of *Superintelligence* in *Quillette* in 2024¹⁶, work as philosophical critique rather than methodological audit. Armstrong's 2012 assessment of Kurzweil is the closest analogue, and the methodology used in the present paper is built on it. Four distinct methods are stacked in Urban's synthesis, and each requires separate examination.

§3.1 The Müller and Bostrom survey method

The 2013 survey targeted four populations: participants of the 2011 PT-AI conference, members of the AGI Conference 2012, members of the editorial board of the journal *Synthese* in its volume on philosophy of artificial intelligence, and the hundred most cited researchers in artificial intelligence according to Microsoft Academic Search. Response rates ranged from 29 to 46 percent depending on the population. The aggregated n was 170. Three structural features of this design matter for the 2026 score.

Self-selection. Three of the four populations are explicitly self-identified as concerned with the long-run trajectory of AI. The AGI Conference attendee population in particular skews strongly toward respondents who expect AGI to be reachable in a tractable timeframe, since attending the conference is itself a partial signal of that expectation. AI Impacts has subsequently estimated that population selection for optimistic respondents accounts for around two decades of acceleration in median timelines compared to broader cohorts¹⁷. The 2013 survey sits at the optimistic end of that range, and the 2022 to 2024 Grace surveys, which use a larger and less self-selected sample of NeurIPS and ICML authors, produce later medians.

Definitional drift. Müller and Bostrom asked about *high-level machine intelligence*, defined operationally as the capacity to carry out most human professions at least as well as a typical human. Urban transmitted this as *artificial general intelligence*, and Kurzweil treats AGI as a stronger and more specific target. By 2026, the relevant benchmark capacity is dispersed across HLE, ARC-AGI-2, SWE-bench Verified, FrontierMath and others, none of which corresponds cleanly to either the survey's HLMI definition or to the colloquial AGI target. A claim that the 2022 ten-percent threshold has been satisfied by GPT-3.5 in November 2022 or by GPT-4 in March 2023, following the Bubeck and colleagues *Sparks of AGI* reading¹⁸, is defensible. A claim that it has not is also defensible.

Aggregation. Müller and Bostrom report aggregated medians rather than full posterior distributions. The reader of the 2013 paper sees that fifty percent of respondents place HLMI before 2040 but does not see the spread of individual judgments or the contribution of population subsets to the aggregate. This compression made the survey citable in popular synthesis at the cost of obscuring the disagreement among respondents.

16. Welsh, S. (2024), *Superintelligence, Ten Years On*, Quillette, 2 July 2024.

17. AI Impacts (n.d.), *Accuracy of AI predictions*, aiimpacts.org.

18. Bubeck, S., et al. (2023), *Sparks of Artificial General Intelligence*, arXiv:2303.12712.

§3.2 The Kurzweil extrapolation method

Kurzweil's method is exponential extrapolation. The Law of Accelerating Returns, in its sharpest form, claims that the doubling time of a wide range of technology metrics is itself shortening, producing super-exponential growth on long time horizons. The compute price-performance series, the genome sequencing cost curve, and the chess Elo trajectory of computers were the favoured exhibits in the 1999 and 2005 books¹⁹. The 2025 milestone for one-thousand-dollar brain-equivalent compute follows from extrapolating the consumer FLOP-per-dollar curve forward from approximately 2003.

Armstrong's 2012 assessment of Kurzweil's 1999 predictions identified three recurring failure modes in this extrapolation. First, the choice of curves was retrospectively favourable to the conclusion: metrics that grew slower or that exhibited S-curve saturation in the 1990s and 2000s are underrepresented in the exhibits. Second, the conversion between physical quantities and target outcomes is loose, in particular the conversion of FLOPS to cognitive capacity. Third, the predictions of *first laboratory demonstration* are systematically more accurate than the predictions of *mass market availability*, but the predictions were typically interpreted as mass-market by readers, including Urban. Armstrong's net assessment of Kurzweil's 1999 predictions was an accuracy of 42 percent on his five-point scale against Kurzweil's self-assessed 86 percent.

What is recoverable from Kurzweil's method is the price-performance series itself, which has continued to behave roughly as he projected at the consumer-device level. What is not recoverable is the inference from physical capacity to cognitive capacity, which the empirical record of 2017 to 2026 has shown to depend on architecture, training procedure and alignment technique to a degree the 2005 framing did not allow for.

§3.3 The Bostrom analytical-philosophical method

Bostrom's contribution is conceptual structure rather than numerical prediction. *Superintelligence* is constructed around thought experiments, decomposition of possibility space into attractor states, and analysis of failure modes of value alignment. The method is analytical-philosophical in the strict sense: it does not produce empirical predictions that can be scored against observation. It produces a vocabulary that organises subsequent thought.

This is the most resistant to retrospective scoring of the three methods, and the most consequential. The Bostromian balance beam continues to organise the conversation in 2026, including this paper. What can be examined is whether the structural decomposition captures the relevant possibilities. Welsh's 2024 critique¹⁶ argues that the balance beam under-weights pragmatic legal stabilisers. The empirical record of 2015 to 2026 supports that critique on one specific point, which is developed in §6.

16. Welsh, S. (2024), *Superintelligence, Ten Years On*, Quillette, 2 July 2024 (cited again).

19. Kurzweil, R. (1999, 2005), *The Age of Spiritual Machines; The Singularity Is Near*, Viking.

§3.4 The Urban synthesis method

Urban's own method is vulgarisation. He compresses three technical inputs into a single accessible narrative aimed at a non-specialist audience. The technique is well executed and the resulting two parts are still read in 2026, an unusual longevity for a popular AI piece. The compression introduces its own distortions, however. Kurzweil's distinction between first-demonstration and mass-market is flattened. The Müller and Bostrom survey distributions are reduced to their medians. The Bostromian balance beam is rendered visually as a literal beam, which over-determines the binary framing. None of these is unique to Urban; they are features of vulgarisation as a genre. They matter here because the 2015 frame as received by general readers is the synthesised version, not the technical originals, and the structural omissions reproduced below are partly artefacts of synthesis rather than failures of the underlying sources.

§3.5 What is recoverable, what is not

From this methodological audit, three elements of the 2015 frame are recoverable under reformulation, and three are not.

Recoverable. The compute trajectory at consumer pricing, anchored on Epoch AI's FLOP-per-dollar series, has continued at roughly the rate Kurzweil projected. The persistence of the capability-versus-safety funding asymmetry is the most robust empirical regularity of the eleven-year window, and the 2015 frame correctly identified it as a feature rather than as a transient. The analytical decomposition of possibility space into attractor states is a useful conceptual tool, on the condition that the enumeration of attractors is open rather than closed.

Not recoverable. The unmediated inference from compute to cognition is incompatible with the empirical record of 2017 to 2026, which has placed architecture, scaling laws and alignment technique as load-bearing intermediates. The planetary uniformity of the compute trajectory is incompatible with the export-control regime that has bifurcated the United States and Chinese trajectories since October 2022. The binary structure of the Bostromian beam, in its strong form claiming exhaustive enumeration of resting states, is incompatible with the regulatory record of the European jurisdiction.

Section 4 reframes the first set. The second set is discarded for cause.

§4. Three reframed hypotheses that survive

The hypotheses identified in §1.2 are now restated. Each is either reformulated to survive the empirical pressure of the 2015 to 2026 record, or discarded with stated reason. This is the section where the paper takes its first analytical position, after three sections of restitution, comparison and methodological audit.

§4.1 H1 reformulated, compute leads to intelligence under transformer, scaling laws, and reinforcement learning from human feedback

The original H1 ran from compute to cognition as if the path between them were unmediated. The empirical record has shown that the path is mediated by three additional ingredients that did not exist in 2015. The transformer architecture, introduced by Vaswani and colleagues in June 2017²⁰, provided the substrate. The empirical scaling laws of Kaplan and colleagues in January 2020²¹, refined by Hoffmann and colleagues for Chinchilla in 2022²², formalised the conditions under which adding compute predictably improves capability. Reinforcement learning from human feedback, operationalised at scale by OpenAI with InstructGPT in early 2022 and then with ChatGPT in November 2022, turned the resulting models into aligned conversational agents whose outputs were usable for the tasks the benchmark cascade measures. None of these was foreseen by the 2015 frame, and the frame did not need to foresee them in detail. What it needed to acknowledge was that compute alone, in the absence of an architecture that scales predictably and an alignment technique that elicits useful behaviour, would not produce the trajectory it predicted.

Figure G4 on the next page shows the trajectory the reformulated H1 actually traces. Training compute at the frontier grows by approximately a factor of five per year on the recent subset, with a coefficient of determination of 0.918 over the 2012 to 2025 window on the longer series. The trend is more rapid than the Moore's law cadence Urban inherited from Kurzweil by a factor of approximately 4.6, which is to say that the source of the acceleration is not lithography. It is capital expenditure on data centres, high-bandwidth memory, networking and parallelisation. The reformulated hypothesis: compute at the frontier grows roughly five-fold per year, conditional on an architecture that absorbs compute productively and an alignment technique that translates raw capability into useful behaviour.

20. Vaswani, A., et al. (2017), *Attention Is All You Need*, NeurIPS 30.

21. Kaplan, J., et al. (2020), *Scaling Laws for Neural Language Models*, arXiv:2001.08361.

22. Hoffmann, J., et al. (2022), *Training Compute-Optimal Large Language Models*, arXiv:2203.15556.




Figure G4. Frontier training compute, 2012 to 2026, projection 2027 to 2030. Log scale. Linear regression on log compute across the AlexNet to Grok 4 series, slope 0.62 OOM per year (equivalent to a factor of 4.2 per year over the longer window), R^2 0.918. The more recent 2020 to 2026 subset gives a faster doubling time of 5.2 months per Epoch AI's February 2026 update. Largest known training run at May 2026 is Grok 4 at approximately 5×10^{26} FLOP. Bottleneck bands annotate the projected activation of three constraints between 2026 and 2030. Sources, Epoch AI Trends, IEA *Energy and AI 2025*, SemiAnalysis production analyses.

The frontier training compute curve thus exhibits a slope of 0.62 orders of magnitude per year over the AlexNet to Grok 4 series, with R^2 of 0.918. The recent 2020 to 2026 subset accelerates further to a doubling time of 5.2 months. The trend is more rapid than the Moore's law cadence by a factor of approximately 4.6, confirming that lithography is no longer the rate-limiting factor at the frontier; capital expenditure on data centres, high-bandwidth memory, networking and parallelisation has displaced it. The bottleneck bands annotated on the figure (HBM supply sold out 2026, United States grid power additions of approximately +130 percent for data centres by 2030 per IEA, capital expenditure ceiling implied by training runs above ten billion United States dollars per run by 2027 to 2028) define the perimeter within which the projection of \$5 is constructed.

§4.2 H4 reformulated and inverted, regulation precedes capability in jurisdictions that legislate ex ante

The original H4 stated that regulation follows technology. The empirical record of 2021 to 2024 has reversed the direction of this relationship in at least one jurisdiction. Regulation (EU) 2024/1689 was tabled by the European Commission in April 2021. At that date, GPT-3 was eleven months old, the transformer-based foundation model paradigm was four years into its empirical existence, and ChatGPT was eighteen months away from public release. The regulation was published in the *Official Journal of the European Union* on 12 July 2024. Its general-purpose AI obligations entered application on 2 August 2025. Its high-risk system enforcement powers activate on 2 August 2026. The Digital Omnibus trilogue agreement of 7 May 2026 redistributes some Annex III dates to 2 December 2027²³ without reversing the underlying architecture.

G5 · AI Act application calendar, post-Omnibus 7 May 2026

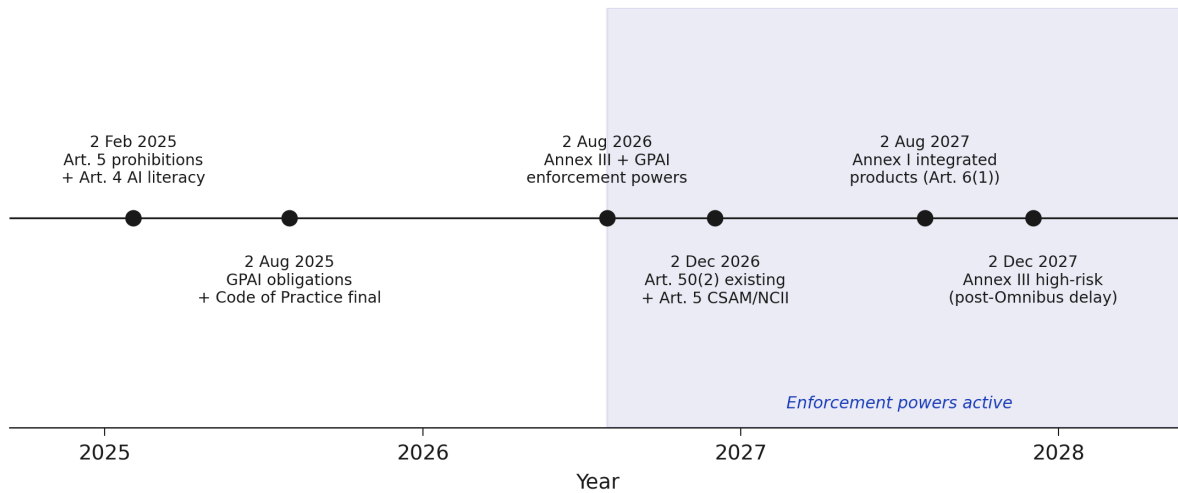


Figure G5. AI Act application calendar after the Omnibus trilogue agreement of 7 May 2026. Six anchor dates. Enforcement powers over GPAI providers activate on 2 August 2026, the same day Annex III high-risk obligations originally applied; the Omnibus moves some of those obligations to 2 December 2027. Article 50(2) watermarking and Article 5 CSAM and NCII safe-harbour provisions apply to systems already on the market on 2 December 2026. Sources, *OJ L* 12 July 2024, European Commission AI Office communications, Digital Omnibus trilogue summary 7 May 2026.

The reformulated H4: in jurisdictions that legislate ex ante, regulation can precede the broad deployment of the technology it constrains, and produce a structural artefact that the technology trajectory must then accommodate. The reformulation is jurisdiction-conditional. The United States revoked Executive Order 14110 in January 2025 and the federal regulatory picture is more fragmented there. The Korean AI Basic Act applies from January 2026 with extraterritorial reach. The United Kingdom is preparing a horizontal frame. But in the European market, which the original H4 implicitly aggregated with all others, the inversion is empirical, not speculative.

23. European Commission and European Parliament (2026), *Digital Omnibus AI Act trilogue agreement*, 7 May 2026.

The reformulation produces a consequence the 2015 frame did not allow for. It introduces a category of public artefact that did not previously exist: the deployer-side, article-by-article documentation cascade required under Article 26 of the regulation, with the registration obligation surviving in streamlined form under Annex VIII Section B of the post-Omnibus consolidated text²³. This category sits between the regulator's certification, which the AI Act does not produce for deployers, and the model itself, which is the responsibility of the provider. It is the object the deployer must produce, on its own initiative, to be in a defensible position. §6 develops the consequences.

§4.3 The asymmetry hypothesis preserved

Urban observed in 2015 that funding for AI safety research lagged funding for AI capability development. He treated this as a transitional condition that the field would notice and correct as capability approached the human threshold. Eleven years on, the asymmetry has not closed; it has widened.




Figure G6. The asymmetry that did not close, 2015 to 2026. Both series on log scale. AI capital expenditure of frontier hyperscalers aggregates the AI-attributable share of disclosed infrastructure spend. Safety funding aggregates the largest publicly identifiable institutional commitments, including the UK AI Security Institute (renamed from UK AI Safety Institute, February 2025), the United States Center for AI Standards and Innovation, CAISI (renamed from US AI Safety Institute, June 2025), the Frontier Red Team at Anthropic, METR, ARC Evals, MIRI, Apollo Research, and identifiable academic alignment grants. Ratio at May 2026 approximately 0.13 percent.

Microsoft's stated calendar-year 2026 capital expenditure on AI infrastructure is approximately 190 billion United States dollars. Combined hyperscaler AI capex for 2026 reaches an estimated 725 billion dollars. The publicly identifiable global budget for AI safety research totals under one billion dollars per year at the most realistic readings. The ratio is approximately 0.13 percent²⁸, narrowed only marginally by partial regulatory transfer under Article 55(1)(d) of Regulation (EU) 2024/1689, which imposes adversarial testing and red-teaming obligations on providers of general-purpose AI models with systemic risk.

The preserved hypothesis: the structural asymmetry between capability investment and safety investment persists across the eleven-year window, the field does not self-correct, and the partial closure mechanisms that have emerged are external to the industry, principally regulatory. This is the most resilient claim in the original frame. It is also the claim Urban himself thought most likely to be transitional, which is the strongest evidence that the frame underestimated the persistence of incentive asymmetries. The dissolution of OpenAI's Superalignment team in May 2024, after publicly committing twenty percent of the company's compute to alignment research, marks the most visible recent moment at which an industry actor failed to maintain the closure trajectory unprompted.

§4.4 H2 discarded, the trajectory is not global and uniform

The 2015 frame treated compute as a planetary quantity. This abstraction was defensible in 2015 and is no longer defensible at May 2026. On 7 October 2022, the United States Bureau of Industry and Security introduced the first major export controls on advanced chips and semiconductor manufacturing equipment to the People's Republic of China²⁴. Subsequent updates on 17 October 2023 and 2 December 2024 expanded the geographic perimeter and introduced China-wide controls on high-bandwidth memory.




Figure G7. The geopolitical fork. Indicative single-chip performance trajectory, frontier-class accelerators, normalised so that the NVIDIA H100 equals 100 on FP16 throughput. United States trajectory follows the A100, H100, H200, B200 and B300 lineage in blue. Chinese trajectory follows the Ascend 910A, 910B and 910C lineage in red, plateauing on SMIC N+1 with HBM constrained. Vertical lines mark BIS actions of October 2022, October 2023 and December 2024.

23. European Commission and European Parliament (2026), *Digital Omnibus AI Act trilogue agreement*, 7 May 2026 (cited again on this page).

24. US Bureau of Industry and Security (2022, 2023, 2024), *Implementation of Additional Export Controls*, Federal Register.

28. The 0.13% ratio aggregates publicly identifiable safety institutional commitments (UK AI Security Institute, US CAISI, MIRI, ARC, METR, Apollo Research) and publicly disclosed safety-focused budgets within frontier labs, notably Anthropic's Frontier Red Team. Internal red-teaming and alignment research investments within OpenAI, Google DeepMind, and Meta, where not separately disclosed, are not included. Including conservative estimates of these undisclosed internal investments would shift the ratio toward 0.2-0.3 percent, without changing the order of magnitude.

The Chinese response, traceable in primary sources, did not collapse frontier capability but bent its trajectory onto a separate curve. Huawei's Ascend 910C began mass shipping in May 2025 on SMIC's N+1 process, equivalent to roughly an older seven-nanometre node without extreme ultraviolet lithography. SemiAnalysis reports the part at approximately 60 percent of the inference performance of an NVIDIA H100²⁵. DeepSeek V3 and R1 were trained on NVIDIA H800. Industry reports through 2025 and 2026 indicate persistent dependence on NVIDIA hardware for frontier-class Chinese training runs, with intermittent transitions to Huawei silicon for specific deployments³². H2 is therefore discarded. The trajectory is no longer global and uniform; it is at least two coupled but distinct trajectories.

§4.5 H3 partially discarded, the binary frame does not exhaust possibility

The Bostromian balance beam contains, in its original framing, two attractor states. §6 will argue that the empirical record of 2015 to 2026 is consistent with a third structural pole, distinct from extinction and from immortality, that the original framing did not contain. Before deferring the doctrinal argument to §6, three reasons for partially discarding H3 are stated here, because they are independent of the regulatory anchoring developed in §6.

First, the binary frame assumes trajectories without stabilising institutional intervention.

Bostrom's balance beam is a mechanical metaphor borrowed from dynamical systems. In the strict sense it borrows from, an attractor is a region of state space toward which trajectories converge in the absence of perturbation. The frame treats the species and the superintelligent artefact as the only candidate sources of perturbation. Law, in the positive-institutional sense, is not modelled as a perturbation source because it is implicitly assumed too slow and too jurisdictional to act on the relevant timescale. The 2015 to 2026 record contradicts that assumption in the European jurisdiction, as §4.2 documents. The frame's exhaustiveness claim is therefore weakened wherever a regulatory regime can act before the relevant capability is broadly deployed.

Second, the binary frame does not distinguish between trajectories toward extinction and trajectories slowed sufficiently for other factors to intervene.

A balance beam is a topology with no third resting state; either the species falls one way or the other. The empirical record of 2015 to 2026 shows that capability deployment can be slowed by regulatory friction in particular jurisdictions, by export-controls cascades, by capital constraints (§5.3 develops this), and by the architectural requirements that the 2015 frame did not anticipate. Each of these introduces time. Time is what other stabilising factors require to act. A frame that compresses the trajectory toward an immediate tip in one direction or the other underweights the structural function of delay.

Third, institutional inertia is itself a force exercising on the trajectory.

The Bostromian frame underweights this systematically because the dynamical-systems metaphor it borrows is mechanical, not institutional. Mechanical systems do not contain bureaucracies, multi-stakeholder consultations, comitology, judicial review, or sector-specific implementing acts. Institutional regimes do, and they generate friction proportional to their complexity. The AI Act has produced, as of May 2026, a Code of Practice with named signatories, an AI Office with a defined remit, a designation procedure under Article 52, a sanctions framework under Article 99, and a cascade of national competent authorities tasked with enforcement. None of these existed in 2014. Each adds friction to whatever trajectory the underlying technology takes. The frame that ignored law in 2014 missed that institutional friction is an active force, not a passive constraint.

H3 is therefore partially discarded: the two-attractor decomposition is preserved as a vocabulary, because it remains useful for organising thought about long-run trajectories. But its claim to exhaust possibility space is rejected. §6 names the third pole and argues for its empirical reality at May 2026.

§5. A constrained minimal path to 2030

This section constructs a minimal projection of the 2026 to 2030 trajectory anchored on the patterns identified in §4. The exercise is not a scenario in the Kokotajlo and colleagues sense²⁶. It does not narrate a month-by-month trajectory. It identifies ranges, anchored on empirically stable rates, with bottleneck activations flagged in their expected order. The IPCC AR6 calibrated language is used for qualitative confidence: *likely* for above 66 percent confidence, *very likely* for above 90 percent, *virtually certain* for above 99 percent²⁷. The language is reserved for cases where the underlying base rate supports it.

§5.1 Compute projection, three trajectories

Anchored on the May 2026 frontier of approximately 5×10^{26} FLOP for the Grok 4 training run, three trajectories are plotted in figure G8. The lower bound assumes a factor of three per year, reflecting bottleneck-constrained growth where high-bandwidth memory supply, grid power additions, fab capacity at N3 and below, and capital balance-sheet ceilings bite earlier than the trend would otherwise predict. The median assumes the Epoch AI factor of five per year, continuing the 2020 to 2026 cadence. The upper bound assumes a factor of seven per year, reflecting accelerated scenarios in which the announced Stargate-class build-outs come online on schedule and supply constraints relax.




Figure G8. Constrained projection of frontier training compute, 2026 to 2030. Three trajectories anchored on the Grok 4 May 2026 figure. Lower bound at $\times 3$ per year, median at $\times 5$ per year, upper bound at $\times 7$ per year. Bottleneck activation markers are placed at HBM supply (mid-2026, sold out at SK Hynix and Micron for 2026), power grid additions (mid-2027, IEA forecast of +130 percent United States data centre demand by 2030), capital expenditure ceiling (2028, training runs above ten billion United States dollars per run). The Article 51 systemic threshold of 10^{25} FLOPs is marked for reference.

Under the median trajectory, the frontier training run reaches approximately 10^{28} FLOP by mid-2028 and 10^{29} FLOP by 2030. Under the lower bound, it reaches approximately 3×10^{27} FLOP by 2028 and 10^{28} FLOP by 2030. Under the upper bound, it reaches 10^{29} FLOP by 2028 and 10^{30} FLOP by 2030. The spread between the bounds is one and a half orders of magnitude at the 2030 horizon, which is the empirical uncertainty band.

§5.2 Capability projection and the benchmark question

The relationship between compute and capability at the frontier is the subject of an active research literature. Empirically, the benchmark cascade of §2 shows that benchmark scores have moved roughly in lockstep with compute across the recent generation, with twelve-to-twenty-four-month saturation cycles. Whether this rate persists into the projection horizon depends on three things. First, whether the benchmarks themselves remain falsifiable as model capability rises; HLE was designed to resist saturation and is at 44.7 percent twelve months after release, which suggests the institutional production of saturation-resistant benchmarks is keeping pace, though not by a wide margin. Second, whether the architectural innovations continue at the recent pace; test-time compute scaling, mixture of experts, and tool use have been the principal sources of capability gain in 2024 to 2026. Third, whether data exhaustion intervenes; Villalobos and colleagues at Epoch AI estimate effective stock of quality and repetition-adjusted human text at approximately 300 trillion tokens, with full utilisation between 2026 and 2032 under continued trends³⁵.

It is *likely*, in the IPCC calibrated sense, that frontier capability scores on benchmark families released after 2026 will reach saturation on at least three of the six families currently in use by 2028, conditional on the median compute trajectory.

§5.3 Capex and bottlenecks, an activation order

The compound annual growth rate of hyperscaler AI capital expenditure between 2020 and 2026 is approximately eighty percent. Extrapolation at this rate is mechanically untenable: hyperscaler balance sheets do not accommodate it past 2028 to 2029 under conventional financing structures. The lower bound of the compute trajectory in figure G8 reflects this. The bottleneck activation order, based on current primary sources, is the following.

- **HBM supply** is sold out for 2026 across SK Hynix and Micron, and tightness extends into 2027. Pricing pressure activates first, with negotiated twenty percent price increases on HBM3E contracts for 2026 already reported.
- **Grid power additions** are the constraint most aggressively flagged by the International Energy Agency in its 2025 *Energy and AI* report²⁹, projecting data centre electricity demand to rise from approximately 485 TWh in 2025 to 950 TWh in 2030, with the United States carrying a 130 percent increase. This constraint binds from late 2026 onward.
- **Fab capacity** at TSMC N3 is reported above one hundred percent utilisation in the second half of 2026 according to SemiAnalysis. N2 capacity is the next-generation bottleneck and is expected to bind through 2028.
- **Capital** is the slowest-moving constraint and the hardest to forecast. Training runs above ten billion United States dollars per run are projected for 2027 to 2028 by Epoch AI, which is the threshold at which non-hyperscaler labs are squeezed out of the frontier on financing grounds alone.

25. SemiAnalysis (2025), *Huawei Ascend Production Ramp*, 8 September 2025.

26. Kokotajlo, D., Lifland, E., Larsen, T., Dean, R., with editorial input from Alexander, S. (2025), *AI 2027*, AI Futures Project, April 2025.

27. IPCC (2021), *Climate Change 2021, The Physical Science Basis*, WG I AR6, Chapter 1.

35. Villalobos, P., et al. (2024), *Will we run out of data?*, Epoch AI.

29. International Energy Agency (2025), *Energy and AI*, [iea.org/reports/energy-and-ai](https://www.iea.org/reports/energy-and-ai).

32. Reuters, citing *The Information*, reported in April 2026 that DeepSeek V4 was expected to run on Huawei Ascend chips, after R2 training on Huawei hardware had been stalled by hardware stability, chip-to-chip interconnect, and Huawei CANN software toolkit limitations, prompting a return to NVIDIA hardware for training while inference shifted to Ascend 910C. Cross-corroboration in Tom's Hardware coverage of the same Reuters/The Information report, April 2026.

§5.4 The regulatory corollary of the projection

The compute projection has a regulatory corollary that the 2015 frame did not allow for. Article 51 of Regulation (EU) 2024/1689 designates a general-purpose AI model as carrying systemic risk when its cumulative compute used for training exceeds 10^{25} FLOPs, with the threshold revisable by the European Commission via implementing act⁵. At May 2026, the European AI Office estimates that between five and fifteen models worldwide currently exceed this threshold³⁰.

Under the median compute projection of §5.1, the count of models at or above the 10^{25} FLOPs threshold scales with the number of frontier training runs and with the dispersion of those runs across providers. Under the assumption that the top five frontier providers each release approximately two training runs above threshold per year, with attrition for discontinued models, the count rises from the AI Office's May 2026 estimate of between five and fifteen models toward an order of magnitude of approximately thirty (± 10) by 2 August 2028 and approximately sixty (± 20) by 2030. These figures are sensitive to three factors: (a) whether the Commission revises the threshold upward by implementing act, (b) whether the count is taken at the model level or at the family level, and (c) whether non-frontier providers cross threshold via fine-tuning of open-weight base models. The order of magnitude is robust to these sensitivities; the precise count is not.

Each such model triggers Article 51 to Article 55 systemic risk obligations on its provider: model evaluation including adversarial testing, serious incident reporting to the AI Office, and cybersecurity protections proportionate to the systemic risk identified. Each such model also has deployers downstream. An empirical screening conducted in April 2026³¹ indicates a near-universal absence of public regulatory positioning across a sample of fifty European AI providers: of the fourteen entities in the sample holding existing CE marks under MDR or IVDR, or banking licences under ACPR or ECB supervision, none had mapped the AI Act layer onto their existing compliance artefacts at that date.

Whether the 96 percent proportion holds at scale across the broader European deployer population remains an open empirical question. The present paper does not attempt to settle it; subsequent screenings, including future revisions of the parent report, will refine the population estimate. What can be projected with the median compute trajectory is that the absolute number of deployers exposed to Article 26 obligations scales mechanically with the proliferation of high-risk systems under Annex III, and that the post-Omnibus Annex VIII Section B registration obligation, which survives the simplification in streamlined form, will produce a public artefact for an increasing population of deployers between 2 August 2026 and 2 December 2027. The volume of documentation that must be produced to be in a defensible position rises with the volume of systems, and it does so on a calendar that the regulation has already fixed.

5. Regulation (EU) 2024/1689, *Official Journal of the European Union*, OJ L, 12 July 2024 (cited again).

30. European Commission, AI Office (2025), *Guidelines for providers of general-purpose AI models*.

31. Shucrani, L. B. (2026), *EU AI Act Readiness Report*, April 2026, Zenodo DOI 10.5281/zenodo.19671329. The 96 percent figure (forty-eight of fifty companies) and the 0/14 sub-finding on certified entities are drawn from Findings 2 and 3 of that report.

Annex A, The Deployer Multiplier (May 2026, Zenodo DOI 10.5281/zenodo.20042175) extends the analysis to three downstream deployer segments.

§5.5 The task horizon trajectory, a duration-based complement to the score projection

The compute and capability projections of §5.1 and §5.2 are based on training-time inputs and on score-based benchmarks. Both are necessary, both have known limitations: compute does not translate one-to-one into deployable capability, and score-based benchmarks saturate. METR (Kwa et al., 2025) introduces a duration-based capability metric that is structurally complementary, in the sense that it cannot saturate by design and is calibrated to human task duration rather than to a model-specific score range. The metric is reproduced in figure G9: the 50 percent task-completion time horizon doubles every 207 days on the 2019 to 2025 frontier sample, with a 95 percent confidence interval of 166 to 240 days and a coefficient of determination of 0.80^{31b}.

If the fitted trajectory is extrapolated naively under the same constraints as the compute projection of §5.1 (no architectural break, sustained capital, no policy-imposed slowdown other than the AI Act calendar already in force), the 50 percent task horizon would reach the order of one working day, eight hours of human-equivalent task, within approximately twenty to twenty-four months from the May 2025 anchor. METR itself flags the methodological limits of such an extrapolation: the metric is benchmarked on software-engineering and reasoning tasks, the human baselines are calibrated within a particular task family, and the 50 percent threshold is one design choice among several (an 80 percent threshold doubles more slowly). The trajectory is included here for two reasons. First, because it provides a metric of the kind §2 found absent from the 2015 frame, where Urban relied on Kurzweilian extrapolation of compute and on benchmark scores that did not yet exist for the most demanding tasks; the METR horizon is the kind of metric that, if it had existed in 2015, would have constrained the 2030 superintelligence claim. Second, because the duration framing has direct operational consequences for Article 14 oversight, developed in §6.5.




Figure G9. METR 50% task-completion time horizon, 2019 to 2025 frontier sample, with confidence band. Each point is a frontier model at its release date, plotted against the duration of human task at which the model achieves a 50 percent success rate on the METR task suite. Fitted exponential trajectory, doubling 207 days; 95 percent confidence interval, 166 to 240 days; R^2 equals 0.80. Source, Kwa et al. (2025), METR, arXiv:2503.14499v3 (February 2026 revision).

31b. Kwa, T., West, B., Becker, J., et al. (2025), *Measuring AI Ability to Complete Long Tasks*, METR, arXiv:2503.14499v3, February 2026 revision (cited at §2.5 fn 13b).

This corollary is not a forecast of demand for any service. It is a mechanical consequence of the compute trajectory and the regulatory architecture. Whether the corollary is recognised by deployers, treated as load-bearing by their boards, and converted into operational documentation by the relevant calendar dates, is a separate question. The corollary itself is structural and is consistent with both the median and lower-bound compute trajectories.

§6. The third structural pole, definition, empirical grounding, operational consequences for European deployers

The previous five sections were retrospective and projective; the empirical record was the protagonist and the author's interventions were limited to scoring and reformulation. What follows is a doctrinal argument that the empirical record is consistent with, but does not by itself entail. The reader is invited to accept the empirical record and reject the doctrinal frame, or to accept both. They are separable.

This paper proposes that the empirical record of 2015 to 2026 is consistent with the emergence of a *third structural pole*, a regulated deployment regime created by Regulation (EU) 2024/1689, that the original predictive frame did not contain. The pole is not an attractor in the strict dynamical-systems sense (it is not absorbing, not irreversible); it is a stable-by-construction institutional layer that any future trajectory must accommodate in the European jurisdiction. The original Bostromian frame (2014), transmitted by Urban (2015), posited two attractor states and treated positive law as too slow, too national, and too reactive to act as anything other than after-the-fact constraint. This paper does not contest the frame as a philosophical exercise; it observes that the frame omitted positive law as a candidate stabiliser, an omission that the empirical record of 2015 to 2026 in the European jurisdiction no longer permits.

§6.1 The empirical case for a third structural pole

The empirical case rests on three observations, all visible in the primary record summarised in §2, §4 and §5. The operational instantiation of the pole is documented in *Annex A*, §8 (the four deployer-side outputs).

First, the regulation precedes the technology it constrains. Regulation (EU) 2024/1689 was tabled in 2021, before the public deployment of GPT-3 at scale, and enters enforcement on 2 August 2026, before the deployment trajectories of frontier general-purpose AI have stabilised. This is the opposite of the Bostromian assumption.

Second, the regulation creates a category of artefact that the predictive frame did not contain. The deployer-side, article-by-article documentation cascade, which the language of Annex VIII Section B post-Omnibus refers to as a *public artefact requiring documented justification*, is a third thing. It is not the system itself, and it is not a regulator's certification. It is what the deployer must produce, on its own initiative, to be in a defensible position before a competent authority or a counterparty's procurement function. The 2015 frame did not anticipate that this category would exist at all, much less become the load-bearing artefact in the European market.

Third, the regulatory architecture has been designed by people who took the Bostromian and Yudkowskian arguments seriously. Gabriele Mazzini, who led the European Commission's drafting team for the AI Act and is now a research associate at MIT Connection Science and a senior consultant at the Centre on Regulation in Europe³³, has stated publicly in interviews from 2024 onward that the regulation operationalises a principle of proportionality in a domain where the risks are partially known and the technical state of the art is moving. In an interview with Bright Eye Venture Capital published in April 2025, Mazzini argued that *a more nuanced framework could have introduced proportionality without sacrificing safety*³⁴, indicating that the architect himself reads the AI Act as a first attempt at the structural problem rather than a final equilibrium. The architecture is reflective on itself, in a way the Bostromian binary does not allow for.

33. Centre on Regulation in Europe (n.d.), *Gabriele Mazzini, biography*, cerre.eu/biographies/gabriele-mazzini.

34. Mazzini, G. (2025), interview, Bright Eye Venture Capital, Part 2, April 2025.

§6.2 What the third structural pole is not

The third structural pole is not a certification, nor the AI Office itself, nor a global regime. It is the documentation cascade described in *Annex A*, §8. Its exportability outside the European Union is an open question not addressed in this paper.

§6.3 The doctrinal claim, stated plainly

BOX 2

The third structural pole stated

Between extinction and immortality, the empirical record of 2015 to 2026 has installed a third structural pole, a regulated deployment regime that did not exist in the 2015 frame and that now structures the conditions of any further trajectory in the European market. This claim is falsifiable; see the constrained prediction in §6.5.

The claim is doctrinal. It is not a measurement. It can be rejected by readers who think the AI Act is a transitional artefact that will be displaced by either capability-driven collapse or by international coordination breakdown. Both of those alternatives remain coherent.

The narrower claim of this paper, on which §6.4 and §6.5 build, is that any predictive frame for AI futures that treats Article 51 to Article 55 of the AI Act, and the Code of Practice they support, as ignorable will be inconsistent with the recorded behaviour of the relevant actors from August 2025 onward. The 2015 frame is one such frame. The 2025 frame of Kokotajlo and colleagues' *AI 2027*²⁶ may be another. The two-branch bifurcation in that paper contains a takeover trajectory and a coordinated-slowdown trajectory; neither branch contains a structural pole of the kind described here.

26. Kokotajlo, D., Lifland, E., Larsen, T., Dean, R., with editorial input from Alexander, S. (2025), *AI 2027*, AI Futures Project, April 2025 (cited again).

§6.4 Operational consequences for European deployers

The compute projection of §5 has direct operational consequences for the population of European deployers, in addition to the consequences for the providers covered by Article 51 to Article 55. Four categories of operational consequence follow from the projection.

Already activable today. Article 4, AI literacy, has been in application since 2 February 2025 and applies to all deployers of AI systems regardless of risk classification. Article 5, the eight prohibited practices, has been in application since 2 February 2025 and applies regardless of deployer status. Article 50 transparency obligations on chatbots, deepfakes and synthetic content apply from 2 August 2026. These are observable today and produce enforceable obligations now.

Premature but imminent. Article 26, the twelve paragraphs of deployer obligations on high-risk systems classified under Annex III, applies from 2 August 2026 for the originally scheduled domains and from 2 December 2027 for the Omnibus-deferred domains. Article 27, the Fundamental Rights Impact Assessment, applies on the same calendar with automatic triggers for credit scoring under Annex III paragraph 5(b) and insurance pricing under paragraph 5(c). Article 50(2), the watermarking obligation on systems already on the market, applies from 2 December 2026 under the Omnibus transitional regime.

Grey zones. The AI Office's implementing acts on Article 6(3), the no-significant-risk exemption that generates an Annex VIII Section B registration obligation under the Omnibus, are expected by Q4 2026. Deployers should monitor them. Subsequent revisions of the General-Purpose AI Code of Practice will also refine the technical state of the art on red-teaming, serious incident reporting, and multi-layered marking of synthetic content.

Open doors. Jurisprudence is emerging in parallel. The Amsterdam District Court ruling of 26 March 2026, ECLI:NL:RBAMS:2026:3106, imposed daily penalties of one hundred thousand euros per defendant entity on xAI, X Corp and XIUC for failure to prevent non-consensual intimate imagery under existing Dutch law and GDPR, citing CJEU *Russmedia C-492/23* as authority on joint platform controllership³⁷. The ruling is not an AI Act enforcement action. It is, however, the strongest signal currently available about how Article 5(1) prohibited practices on non-consensual intimate imagery, which apply on 2 December 2026 under the Omnibus, will be enforced by national courts using the apparatus of existing tort and data-protection law.

Direct mapping to Annex A outputs. For each of the three deployer segments examined in the companion *Annex A, The Deployer Multiplier* (HR-tech, MedDev/clinical software, retail banking), the operational consequences are detailed in *Annex A*, §4, §5 and §6 respectively. The present paper does not repeat them; it refers to them as the empirical instantiation of the third structural pole. Any deployer that produces the four outputs defined in *Annex A*, §8 (Article 26 paragraph map, FRIA scope memo, Article 25 boundary log, Article 49 registration) has, by definition, entered the regulated regime that this paper describes.

37. Rechtbank Amsterdam (2026), ECLI:NL:RBAMS:2026:3106, 26 March 2026.

The Article 14 oversight envelope. The METR task horizon noted in §2.5 and §5.5, doubling every 207 days on the 2019 to 2025 frontier sample, has a direct operational consequence for Article 14 of the AI Act. Article 14 requires that high-risk systems be designed and developed in such a way that they can be effectively overseen by natural persons during the period in which they are in use, including with appropriate human-machine interface tools. Where the task horizon of the underlying model approaches or exceeds the duration of an unsupervised execution window, the design of the human-machine interface and the definition of "effective oversight" can no longer rely on continuous human attention as a default assumption. The deployer obligation under Article 26(2) to assign oversight to natural persons with the necessary competence, training and authority, and the provider obligation under Article 14(4) to enable the persons assigned to oversight to correctly interpret the system's output and intervene, must therefore both be reread against a model class whose unit of autonomous action is measured in tens of minutes today and, on the fitted trajectory, in single-digit hours within twenty-four months. The doctrinal point is not that long task horizons defeat Article 14, but that they shift its centre of gravity from "the human reviews each output" to "the human reviews the design of the intervention envelope".

This mapping is the operational reading the present author has produced in collaboration with European deployers and counsel since 2025. The mapping is reproducible by any reader with access to *OJ L* of 12 July 2024 and to the post-Omnibus consolidated text once published.

§6.5 A constrained prediction

The paper closes with one prediction, formulated in the manner of an applied research paper. The horizon is short by the standards of AI forecasting, eighteen months from publication, which is deliberate. The prediction commits the doctrinal frame proposed in §6.1 to §6.3 on the dimension where the Bostromian frame is most sceptical of regulation, the dimension of effective enforcement. If the prediction is wrong, the doctrinal frame loses substantial weight. If it is right, the frame acquires empirical content beyond the present paper.

BOX 3 · CONSTRAINED PREDICTION

Article 99 enforcement on systemic-risk GPAI providers by 2 December 2027

By 2 December 2027, at least one Member State competent authority or the European AI Office will have opened a formal enforcement action under Article 99 of Regulation (EU) 2024/1689 against a designated provider of a general-purpose AI model with systemic risk, for non-compliance with Article 55(1)(d) on adversarial testing and red-teaming obligations.

Horizon. 2 December 2027, eighteen months and ten days from the date of this paper. The horizon is calibrated to the GPAI enforcement window opened on 2 August 2026, with sixteen months of operational margin for a competent authority or the AI Office to investigate, formally open, and publicly disclose a first action.

Conditions of falsification. The prediction is falsified if, at the close of business on 2 December 2027, no public record of a formal Article 99 enforcement action against a designated systemic-risk GPAI provider is available through (a) the AI Office's communications and public register, (b) any Member State competent authority's enforcement disclosure, or (c) media reporting traceable to a verifiable regulatory source. *Opening* the action is sufficient; the prediction does not require completion, sanction, or admission. *Designated systemic-risk GPAI* means a model formally identified by the AI Office under Article 52, or self-notified as exceeding the 10^{25} FLOPs threshold under Article 51 and not contested by the office. *Non-compliance with Article 55(1)(d)* means a finding by the authority that adversarial testing or red-teaming obligations have not been met to the standard the Code of Practice and successor implementing acts will have specified by that date.

Empirical observable. The observable is a public document, dated on or before 2 December 2027, issued by the AI Office or a Member State competent authority, opening an enforcement action against a named provider on the Article 55(1)(d) ground. The observable is verifiable through (a) digital-strategy.ec.europa.eu and the AI Office's official communications, (b) national competent authority enforcement registers, principally those of the data protection and market surveillance authorities designated under the AI Act, and (c) the AI Office's annual report under Article 65 of the regulation.

Why this prediction commits the doctrinal frame. The Bostromian frame underweights the structural role of positive law because law is assumed to act after the fact, slowly, and through coordination problems. The Article 99 enforcement window is the dimension on which this assumption is most directly testable. If no action is taken within sixteen months of the enforcement window opening, the regulatory pole proposed as the third structural pole loses operational weight. If an action is taken, the third structural pole acquires a documented enforcement track record that no predictive frame of the AI domain can ignore.

Confidence. We assess this prediction as carrying around even odds, with substantial weight on either outcome. The historical base rate for first formal enforcement actions under comparable EU regulatory regimes (GDPR, DMA, DSA) sits above eighteen months: the first major Article 58 corrective measure against a frontier technology actor under the GDPR took approximately twenty-two months from 25 May 2018. The prediction therefore commits to a faster-than-base-rate enforcement trajectory. Falsification would not invalidate the doctrinal frame of §6, but would constitute a substantial signal that the regulatory pole remains declaratory rather than operational in its first enforcement cycle. The point of the prediction is not to assert that it will be right; it is to commit the doctrinal frame to a verifiable test.

§6.6 Relationship to the parent report and Annex A

This paper is a companion to the *EU AI Act Readiness Report* (Sprinkling Act, April 2026, Zenodo DOI 10.5281/zenodo.19671329) and its *Annex A, The Deployer Multiplier* (May 2026, Zenodo DOI 10.5281/zenodo.20042175). The empirical findings on which §5.4 and §6.4 of the present paper rest, the classification gap, the deployer multiplier, and the CE-to-AI-Act gap, are drawn directly from those documents. The third structural pole proposed here is not a separate theory; it is a conceptual reframing of the obligations documented in the parent report. Any deployer that produces the four outputs defined in *Annex A*, §8 (Article 26 paragraph map, FRIA scope memo, Article 25 boundary log, Article 49 registration) has, by definition, entered the regulated regime that this paper describes. Future empirical work may extend the analysis to insurance pricing under Annex III paragraph 5(c), to public administration under Annex III paragraph 6, and to LegalTech tools assisting legal professionals; this paper does not anticipate the scope or timing of that work.

§7. Methodology note

This note is generalisable to any retrospective exercise of comparable scope. It is intended to be citable on its own and is the only part of the paper that bears no Sprinkling Act specifics.

§7.1 Atomisation

Predictions are decomposed to atomic claims, indexed C1 through C n . Each claim is annotated with date of issuance, source, falsifiability rating (high, medium, nil), ambition level (conservative, median, audacious), and horizon (settled, near, far). Composite claims are split into their atomic components before scoring. Where a source bundles multiple predictions into a single statement, the bundle is unbundled in the verdicts table, with cross-references between components preserved.

§7.2 Scoring scale

Five levels, after Armstrong (2012)⁶ and Tetlock (2015)³⁸, are used: verified, partially verified, indeterminate, refuted, premature. A claim with non-zero falsifiability that lies inside its horizon receives one of the first four. A claim outside its horizon receives premature and is excluded from the aggregate score until time passes. Non-falsifiable claims are excluded from scoring and noted as such.

§7.3 Sourcing hierarchy

Sources are hierarchically constrained. Level one is quantitative primary: Epoch AI Trends, MLPerf, Stanford AI Index, 10-K and 10-Q filings, IEA, BIS Federal Register. Level two is primary research: arXiv with citation depth above one hundred, lab research blogs, AI Impacts surveys, the Existential Risk Persuasion Tournament. Level three is regulatory institutional: *OJ L*, AI Office communications, ENISA, EDPB, NIST. Level four is economic data, used only when corroborated by level one or two. Op-ed sources are excluded as sole basis.

§7.4 Magnitude versus direction

Each claim is scored on direction (correct, incorrect) and on magnitude (within factor 2, within factor 10, beyond order of magnitude) separately. A claim can be directionally correct and magnitudinally wrong, in which case it receives partially verified with a magnitude flag. This is the most common verdict in the Urban corpus.

6. Armstrong, S. (2012), *Assessing Kurzweil, the results*, LessWrong (cited again).

38. Tetlock, P. E., and Gardner, D. (2015), *Superforecasting, The Art and Science of Prediction*, Crown.

§7.5 Regulatory overlay

Each claim is mapped to the article or articles of the relevant regulation it would activate if realised. For AI predictions this is principally Regulation (EU) 2024/1689, Articles 5, 6, 26, 27, 50, 51 through 55, and 99. The overlay does not score the regulation against the prediction. It identifies the obligations that would have been triggered, which is a separate epistemic operation from the empirical scoring.

§7.6 Uncertainty language

Following the IPCC AR6 framing, qualitative confidence is reported as *likely* (above 66 percent), *very likely* (above 90 percent), *virtually certain* (above 99 percent), and their negations. The language is reserved for cases where probabilistic estimation is defensible and is not used as decoration.

§7.7 What this methodology does not produce

It does not produce a single accuracy number. Aggregate scores are reported as distributions over the five levels, not as percentages. The temptation to compress a directionally-correct-magnitudinally-wrong claim into a single percentage is the chief cause of the Kurzweil and Armstrong disagreement of 2012, and the methodology resists it by construction. A reader who needs a single accuracy figure for institutional reasons is encouraged to compute one from the distribution; the construction of that figure is a separate methodological exercise that the present paper does not endorse but does not prevent.

The methodology also does not produce a prediction. The constrained prediction in §6.5 is generated from the doctrinal frame of §6 and from the projection of §5, not from the scoring methodology itself. The scoring methodology is intended to evaluate predictions, not to produce them. Confusing the two functions produces a class of retrospective papers that score themselves on the predictions they produce, which is a recognised failure mode of the genre.

§7.8 Reclassification note on C7

C7 (Kurzweil's claim that by 2021 as much progress would be made as in the entire 20th century) is reclassified from indeterminate to non-falsifiable, alongside C16-C19, on the grounds that the metric "progress equivalent to a century" is not operationalised at the source level. Maintaining indeterminate as a separate verdict category for a claim that cannot be operationalised would dilute the methodological distinction between insufficient data and unmeasurable predicate. The reclassification preserves the methodological consistency of the scoring.

§7.9 Scope of this note

The note above is the only part of the paper that bears no Sprinkling Act specifics. It is reproduced here in expectation that future retrospective exercises on comparable predictive frames, whether on Bostrom 2014, Kurzweil 2005, Kokotajlo 2025, or successors, will benefit from a shared methodological vocabulary. Citation: §7 of the present paper, with the paper's full bibliographic reference.

Appendix A. Verdicts table, twenty-five atomic claims

The table below scores each atomic claim extracted from Urban (2015) and its upstream sources. Non-falsifiable claims are excluded from the aggregate distribution shown in figure G2. Premature claims are flagged but not assigned a substantive verdict.

Ci	Claim (condensed)	Verdict	Magnitude or note
C1	Median expert AGI = 2040 (Müller and Bostrom 2013)	PARTIAL	Grace 2024 places median HLMI 50% at 2047; direction correct, late by approximately 7 years
C2	10% AGI by 2022, 50% by 2040, 90% by 2075	PARTIAL	2022 borderline satisfied under Bubeck "Sparks of AGI" reading; definitional drift acknowledged
C3	Kurzweil AGI = 2029	PREMATURE	Within horizon; signals favourable (ARC-AGI-2 85% May 2026, HLE 44.7%)
C4	Kurzweil ASI / singularity = 2045	PREMATURE	T+19, not scored
C5	\$1,000 PC reaches 10 ¹⁶ cps by 2025	PARTIAL	FLOP/\$ trajectory holds; brain-cps conversion contested between 10 ¹⁴ and 10 ¹⁷ FLOPS
C6	Whole-brain reverse engineering ≈ 2030	REFUTED	FlyWire 139,255 neurons (2024); MICrONS 200,000 mouse cortex (2025); human gap remains ~6 OOM
C7	By 2021, as much progress as the 20th century	NON-FALSIFIABLE	Metric not operationalised at the source level. Reclassified alongside C16-C19 as non-falsifiable rather than indeterminate, to maintain methodological consistency (see §7.8)
C8	21st century = 1,000× progress of 20th	PREMATURE	T+74, not scored
C9	Moore's Law continues (doubling ~2 years)	PARTIAL	Wrong mechanism (transistor density slowed since 2016); right conclusion (compute doubles); magnitude faster (5.2 months at the frontier)
C10	AGI to ASI rapid (≤2 years): 10% probability	PREMATURE	AGI not yet unambiguous; not scored
C11	AGI to ASI ≤30 years: 75% probability	PREMATURE	Horizon not yet reached
C12	ASI ≈ 2060 (Urban synthesis)	PREMATURE	T+34, not scored
C13	Outcomes 52% good / 31% bad / 17% neutral (Müller-Bostrom)	REFUTED	Distribution moved toward pessimism: XPT (Karger 2024) 6% experts / 1% superforecasters; Grace 2024 38-51% assign ≥10% to extremely bad
C14	Barrat AGI Conference: 42% AGI by 2030	PARTIAL	Order of magnitude held; selection bias acknowledged
C15	Recursive self-improvement as primary mechanism	PARTIAL	Emerges as a layer (LADDER, AlphaEvolve, Apollo scheming evals) not as primary architecture

Table continues on page 32. Aggregate distribution and reading note follow. C7 reclassified from indeterminate to non-falsifiable for methodological consistency (see §7.8).

Appendix A (continued). Verdicts table, C16 to C25

Ci	Claim (condensed)	Verdict	Magnitude or note
C16-C19	Singleton; incontrollability; goal-spec dystopia; Fermi paradox	EXCLUDED	Non-falsifiable at the current stage; reasoning excluded from the aggregate score
C20	Hardware speed 10 ⁷ × brain	PARTIAL	Raw frequency correct; energy/op comparison reverses the asymmetry (brain ~2-3 OOM more efficient)
C21	Safety funding lags capability funding	VERIFIED	Ratio at May 2026 ≈ 0.13%; partial regulatory closure via Article 55(1)(d)
C22	Three paths to AGI: emulation / GA / self-improving	REFUTED	Effective path is transformer + scaling + RLHF; absent from the frame
C23	ANI examples 2015 (Watson, Siri, search ranking)	VERIFIED	Trivially verified
C24	Algorithmic trading > 50% volume US	VERIFIED	JPM 2023 60-73%; HFT market-making 72.3% (2024); claim became conservative
C25	WBE state of art 2015 = flatworm 302 neurons (OpenWorm)	VERIFIED	Correct in 2015; trajectory advanced ~3 OOM by 2024 (FlyWire)

Aggregate distribution

Across the twenty-five atomic claims, the distribution is the following: four verified (C21, C23, C24, C25), seven partially verified (C1, C2, C5, C9, C14, C15, C20), three refuted (C6, C13, C22), six premature (C3, C4, C8, C10, C11, C12), and five non-falsifiable (C7 reclassified from indeterminate; C16-C19 excluded for non-falsifiability of singleton/incontrollability/goal-spec dystopia/Fermi paradox). The aggregate signal is directional correctness with systematic magnitude bias.

The four claims that are unambiguously verified concentrate on retrospective factual observation (the state of narrow AI in 2015, the dominance of algorithmic trading, the OpenWorm baseline) and on the structural asymmetry between capability and safety funding. The three refuted claims concentrate on architectural prediction (whole-brain reverse engineering by 2030, three paths to AGI) and on outcome distribution among experts (the 2013 optimism that has not survived the empirical record). The seven partial verifications concentrate on timelines and on hardware-cognition inferences, where direction held and magnitude shifted. The pattern is consistent with the methodological audit of §3.

Reading note. The exclusion of C16 through C19 from the aggregate score does not mean the underlying questions are unimportant. It means that an empirical scoring methodology cannot adjudicate them at the present stage. A philosophical or normative treatment of those questions is a separate exercise.

Appendix B. Acknowledgements, conflict of interest, references

Acknowledgements

This paper was developed independently. No funding was received from any organization with a stake in the regulatory outcomes discussed. The author thanks the maintainers of Epoch AI Trends, the Stanford Institute for Human-Centered AI for the AI Index reports, METR for the task horizon dataset (Kwa et al., 2025), and the authors of MDCG 2025-6 / AIB 2025-1, whose primary documentation made article-by-article cross-references practicable. Errors of fact or interpretation remain the author's alone.

Conflict of interest disclosure

The author operates Sprinkling Act, a pre-conformity advisory firm in the EU AI Act compliance market. Sprinkling Act's commercial activity is described in the companion *EU AI Act Readiness Report* (April 2026, DOI 10.5281/zenodo.19671329). At the time of this paper's publication, no commercial relationship exists between Sprinkling Act and the entities, providers, or sources cited in this paper. The doctrinal argument advanced in §6 is consistent with the author's commercial activity. Readers may reasonably weigh this alignment when assessing the doctrinal claim. The empirical findings in §2-§5, sourced exclusively to external primary documentation, are independent of the doctrinal claim and can be accepted or rejected separately, as stated in §6.4.

HOW TO CITE THIS PAPER

APA: Shucrani, L. B. (2026). *The AI Act as a Third Structural Pole. A Constrained Retrospective on the Predictive Frame (2015 to 2026)*. Sprinkling Act Discussion Paper DP-2026-001. Zenodo. <https://doi.org/10.5281/zenodo.20343244>

BibTeX:

```
@techreport{shucrani2026pole,
  author = {Shucrani, Lamar B.},
  title = {The AI Act as a Third Structural Pole.
    A Constrained Retrospective on the
    Predictive Frame (2015 to 2026)},
  institution = {Sprinkling Act},
  type = {Discussion Paper},
  number = {DP-2026-001},
  year = {2026},
  month = {May},
  doi = {10.5281/zenodo.20343244},
  url = {https://doi.org/10.5281/zenodo.20343244}
}
```

Recommended reading

- **Primary regulatory sources:** *OJ L* 2024/1689 (12 July 2024); Code of Practice for General-Purpose AI, AI Office (10 July 2025); MDCG 2025-6 / AIB 2025-1.
- **Empirical foundations:** Epoch AI Trends (epoch.ai/trends); Stanford AI Index 2025; Grace et al. (2024), arXiv:2401.02843; Kwa et al. (2025), METR, arXiv:2503.14499.
- **Methodological references:** Armstrong (2012); Tetlock and Gardner (2015); IPCC AR6 WG I, Chapter 1.
- **Companion Sprinkling Act publications:** *EU AI Act Readiness Report* (April 2026, Zenodo DOI 10.5281/zenodo.19671329); *Annex A, The Deployer Multiplier* (May 2026, Zenodo DOI 10.5281/zenodo.20042175).

About the author

Lamar B. Shucrani is the founder of Sprinkling Act (Brussels, BCE BE 1034.962.482), a pre-conformity advisory firm focused on the deployer-side cascade under Regulation (EU) 2024/1689. The work focuses exclusively on the deployer-side documentation cascade: use case classification, Article 26 paragraph mapping, Article 27 FRIA scoping where applicable, and Article 25 boundary documentation. Two prior publications underpin the present discussion paper: the *EU AI Act Readiness Report*, April 2026, Zenodo DOI 10.5281/zenodo.19671329; and its *Annex A, The Deployer Multiplier*, May 2026, Zenodo DOI 10.5281/zenodo.20042175. ORCID 0009-0002-5093-8550.

References

1. Armstrong, S. (2012). *Assessing Kurzweil, the results*. LessWrong.
2. Bostrom, N. (2014). *Superintelligence, Paths, Dangers, Strategies*. Oxford University Press.
3. Bubeck, S., et al. (2023). *Sparks of Artificial General Intelligence*. arXiv:2303.12712.
4. Centre on Regulation in Europe (n.d.). *Gabriele Mazzini, biography*. cerre.eu/biographies/gabriele-mazzini.
5. Dorkenwald, S., et al. (2024). *Neuronal wiring diagram of an adult brain*. *Nature*, 30 October 2024.
6. Epoch AI (2026). *Trends in Artificial Intelligence*, update of 5 February 2026. epoch.ai/trends.
7. European Commission and European Parliament (2026). *Digital Omnibus AI Act trilogue agreement*, 7 May 2026.
8. European Commission, AI Office (2025). *General-Purpose AI Code of Practice, Final Version*, published 10 July 2025.
9. European Commission, AI Office (2025). *Guidelines for providers of general-purpose AI models*.
10. Grace, K., Stewart, H., Sandkühler, J. F., Thomas, S., Weinstein-Raun, B., Brauner, J. (2024). *Thousands of AI Authors on the Future of AI*. arXiv:2401.02843.
11. Hoffmann, J., et al. (2022). *Training Compute-Optimal Large Language Models*. arXiv:2203.15556.
12. International Energy Agency (2025). *Energy and AI*. iea.org/reports/energy-and-ai.
13. IPCC (2021). *Climate Change 2021, The Physical Science Basis*. WG I AR6, Chapter 1.
14. Kaplan, J., et al. (2020). *Scaling Laws for Neural Language Models*. arXiv:2001.08361.
15. Kaplan, J., for Meta Platforms (2025). *Statement on the EU Code of Practice for General-Purpose AI*, 18 July 2025.
16. Karger, E., Rosenberg, J., Jacobs, Z., Hickman, M., Tetlock, P. E. (2025). *Subjective-probability forecasts of existential risk*. *International Journal of Forecasting*, vol. 41, issue 2.
17. Kokotajlo, D., Lifland, E., Larsen, T., Dean, R., with editorial input from Alexander, S. (2025). *AI 2027*. AI Futures Project, April 2025.
18. Kurzweil, R. (1999, 2005). *The Age of Spiritual Machines* and *The Singularity Is Near*. Viking.

References (continued)

19. Mazzini, G. (2025). Interview, Bright Eye Venture Capital, Part 2, April 2025.
20. MICrONS Consortium (2025). *Functional connectomics spanning multiple areas of mouse visual cortex*. Nature, 2025.
21. Müller, V. C., and Bostrom, N. (2014/2016). *Future progress in artificial intelligence, a survey of expert opinion*. Springer.
22. Rechtbank Amsterdam (2026). ECLI:NL:RBAMS:2026:3106, 26 March 2026.
23. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 (Artificial Intelligence Act). *Official Journal of the European Union*, OJ L, 12 July 2024.
24. Sandberg, A., and Bostrom, N. (2008). *Whole Brain Emulation, a Roadmap*. FHI Technical Report 2008-3.
25. SemiAnalysis (2025). *Huawei Ascend Production Ramp, HBM is the Bottleneck*. 8 September 2025.
26. Shucrani, L. B. (2026). *EU AI Act Readiness Report*. Sprinkling Act, April 2026. Zenodo, DOI 10.5281/zenodo.19671329.
27. Shucrani, L. B. (2026). *Annex A, The Deployer Multiplier*. Sprinkling Act, May 2026. Zenodo, DOI 10.5281/zenodo.20042175.
28. Tetlock, P. E., and Gardner, D. (2015). *Superforecasting, The Art and Science of Prediction*. Crown.
29. Urban, T. (2015). *The AI Revolution, Parts 1 and 2*. Wait But Why, 22 and 27 January 2015.
30. US Bureau of Industry and Security (2022, 2023, 2024). *Implementation of Additional Export Controls*. Federal Register, 7 October 2022, 17 October 2023, 2 December 2024.
31. Vaswani, A., et al. (2017). *Attention Is All You Need*. NeurIPS 30.
32. Villalobos, P., et al. (2024). *Will we run out of data?* Epoch AI.
33. Welsh, S. (2024). *Superintelligence, Ten Years On*. Quillette, 2 July 2024.
34. Yudkowsky, E. (2008). *Artificial Intelligence as a Positive and Negative Factor in Global Risk*. In Bostrom and Cirkovic (eds.), Oxford University Press.

Appendix C. Limitations and disclaimers

This discussion paper is an analytical reference, not a peer-reviewed retrospective and not an audit certified by a third party. The reader who relies on it for operational, legal, or investment decisions should be aware of the following limitations.

Doctrinal versus empirical. §6 advances a doctrinal argument that is consistent with, but not entailed by, the empirical record presented in §2 through §5. A reader may accept the empirical findings while rejecting the doctrinal frame. The two are separable, as stated in §6.3.

The constrained prediction is a test, not advice. The constrained prediction in §6.5 is falsifiable by construction and is intended to commit the doctrinal frame to a verifiable outcome by 2 December 2027. It is not a recommendation, legal advice, investment guidance, or a forecast of any specific provider's behaviour. The author assesses the prediction as carrying around even odds.

Not a substitute for the operational documents. This paper is a doctrinal companion to the *EU AI Act Readiness Report* (April 2026) and its *Annex A, The Deployer Multiplier* (May 2026). A reader seeking operational guidance on deployer-side obligations, the four outputs of *Annex A* §8, or specific segment analyses (HR, healthcare, retail banking) should start with those documents.

Sprinkling Act is not a regulator. Sprinkling Act is not a law firm, not a Notified Body, not a certification body, and is not affiliated with the European Commission, the European Parliament, the AI Office, the Artificial Intelligence Board, or any national supervisory authority. Organisations recognising themselves in any of the segments developed in the companion *Annex A* should consult qualified legal counsel and, where applicable, the relevant sector supervisor or Notified Body before making compliance decisions.

Author independence with disclosed alignment. The doctrinal argument advanced in §6 is consistent with the author's commercial activity at Sprinkling Act, as disclosed in Appendix B. The empirical findings in §2 through §5 are sourced exclusively to external primary documentation and can be accepted or rejected independently of the doctrinal claim.

Source and method transparency. The retrospective scoring methodology follows Armstrong (2012) and Tetlock (2015). The five-level scale, the atomisation of claims to C1 through C25, the sourcing hierarchy, and the magnitude-versus-direction separation are documented in §7 and are reproducible by any reader with access to the primary sources cited in Appendix B. The methodology note in §7 is intended to be citable on its own.

Horizon of the projection. The constrained minimal path to 2030 (§5) is a range, anchored on empirically stable rates with bottleneck activations flagged in their expected order. It is not a scenario in the Kokotajlo and colleagues sense. The upper and lower bounds reflect documented uncertainty in capital expenditure, grid power, fab capacity, and high-bandwidth memory supply, not subjective confidence.

Errata policy. Material corrections are reported transparently via versioned re-publication (v1.1, v1.2, ...) with an explicit change log. The OpenTimestamps proof of each edition remains valid for that edition. See Appendix D for the cryptographic timestamp details.

Appendix D. Document metadata

Field	
Paper version	1.0 · first edition
Drafting period	10 to 22 May 2026 (window opened after publication of the parent's Annex A on 6 May 2026)
Publication date	22 May 2026
Paper DOI	10.5281/zenodo.20343244
Parent document	<i>EU AI Act Readiness Report</i> (Sprinkling Act, April 12, 2026)
Parent DOI	10.5281/zenodo.19671329
Annex document	<i>EU AI Act Readiness Report · Annex A: The Deployer Multiplier</i> (Sprinkling Act, May 6, 2026)
Annex DOI	10.5281/zenodo.20042175
License	Creative Commons Attribution 4.0 International (CC-BY 4.0) · free to download, share, and cite with attribution
Suggested citation	Shucrani, L. B. (2026). <i>The AI Act as a Third Structural Pole. A Constrained Retrospective on the Predictive Frame (2015 to 2026)</i> . Sprinkling Act Discussion Paper DP-2026-001. Zenodo. https://doi.org/10.5281/zenodo.20343244
Author	Lamar B. Shucrani, Founder, Sprinkling Act
Author ORCID	0009-0002-5093-8550
Badge code	SA-DP-20260522-0001 · verifiable at sprinklingact.com/verify/SA-DP-20260522-0001
Resource type	Publication · Working paper / Discussion paper
Language	English
Subject classification	K24 (Cyber Law) · O33 (Technological Change) · O38 (Government Policy)
Keywords	EU AI Act; Regulation (EU) 2024/1689; deployer cascade; Article 26; Article 27 FRIA; Article 25; Article 49; Article 99 enforcement; predictive frame; Bostrom Superintelligence; Kurzweil singularity; Müller-Bostrom survey; retrospective forecasting; Armstrong methodology; Tetlock superforecasting; structural pole; constrained prediction; Digital Omnibus 2026
Cryptographic timestamp	OpenTimestamps detached proof (.ots) anchored to Bitcoin via four calendar servers (a/b OpenTimestamps pools, Eternity Wall, Catallaxy). Status at publication: pending. Mirror: github.com/sprinkling-act/timestamps
Errata and corrections	Reported transparently via versioned re-publication (v1.1, v1.2, ...) with change log; the .ots proof of each edition remains valid for that edition
Independence	No commercial relationship existed, at the time of publication, between Sprinkling Act and any entity, provider, or source cited in this paper. The doctrinal argument in §6 is consistent with the author's commercial activity as disclosed in Appendix B.
Correspondence	contact@sprinklingact.com

The predictive frame of 2015 was directionally correct on asymmetries but structurally blind to the possibility of a third pole, a regulated deployment regime. That regime now exists in the European market. It is not a theory; it is the set of obligations described in Articles 26, 27 and 50 of the AI Act, operationalised in the four outputs of *Annex A*. Any deployer producing those outputs has, by definition, entered the third pole. Any provider failing to map its position remains exposed. The constrained prediction in §6.5 is one verifiable test of the pole's operational weight; the next eighteen to twenty-seven months will resolve it.

The 2015 map was wrong about timelines and right about asymmetries.

What it did not see was that the future would be neither the Bostromian dystopia nor the Kurzweilian apotheosis, but something slower, regulated, and currently in force.

The methodology note above is generalisable.

The reference frame below is verifiable.

The remainder is empirical, and will be tested by the trajectory of the next eighteen to twenty-seven months.

This discussion paper is free to download, share, and reference with attribution.