

# AI / VR — Eight Decades of Cycles

---

*Eighty-three years of summers, winters, and the long quiet rebrands — 1943 to 2026.*

Dr. George Papagiannakis · ORamaVR · FORTH-ICS · University of Crete

**Two technologies, one shared century.** Artificial intelligence and virtual reality have each gone through repeating cycles of hype, breakthrough, disillusionment, and quiet progress — but the cycles are driven by different physics. AI cycles are funding-and-credibility cycles around algorithmic claims. VR cycles are hardware-and-adoption cycles around displays, optics, and silicon. The two fields have synchronised twice (the shared winter of the 1990s, the shared ignition of 2012) and anti-correlated once decisively (the 2022–24 AI eclipse of the metaverse).

Each timeline below reads left-to-right. Marker dots above the centre line are events that pushed the field toward a summer; markers below indicate winter pressure. Coloured period-boundary years appear along the bottom axis: gold for foundations, rust orange for summer, slate blue for winter, grey for quiet years. An appendix at the end of this document gives the full data table of all forty-two timeline events, each with the explanation that appears in the interactive HTML version.

# One foundation period, three summers, two winters, and a long quiet rebrand.

*From McCulloch & Pitts (1943) to ChatGPT (2022): the funding cycles, breakthroughs, and books that ended eras.*




Figure 1 — AI cycles 1943–2026. Foundations (gold) · Summer I 1956–1974 · Winter I 1974–1980 · Summer II 1980–1987 · Winter II 1987–1993 · Quiet rebrand 1993–2012 · Summer III 2012–present.

**Before AI had a name, it had a mathematics.** McCulloch & Pitts (1943) proved that networks of binary threshold neurons could in principle compute any logical function — the seed of every later neural model from Rosenblatt's Perceptron through AlexNet to Transformers. Hebb (1949) supplied the first credible learning rule. Turing (1950) framed the inheritance question. By the time McCarthy convened Dartmouth in 1956, the conceptual scaffolding already existed; what Dartmouth added was a name, a programme, and an organised funding pitch.

**Three summers, two real winters.** The first AI summer was a funding-and-promise bubble around symbolic reasoning that collapsed under combinatorial explosion and overpromising — Lighthill in 1973 and DARPA cutting speech-understanding in 1974 closed the era. The second was an industrial bubble around expert systems and Japan's Fifth-Generation programme, killed when LISP-machine vendors were obsoleted by general-purpose Sun and Macintosh workstations after 1987.

**The interregnum mattered.** The 1993–2011 “quiet” years are sometimes mis-remembered as a winter, but capability advanced steadily under disguised names — machine learning, statistical methods, kernel methods, data mining. Hinton's 2006 deep-belief-network paper rebranded neural nets as “deep learning,” and ImageNet (2009) gave the field a benchmark to chase. AlexNet (2012) was the ignition event; ChatGPT (2022) the public arrival.

# Three waves, two failures of the silicon to deliver what the marketing promised.

*From Heilig's Sensorama (1962) to Apple Vision Pro (2024): hardware cycles bounded by Moore's Law and display physics.*




Figure 2 — VR / metaverse cycles 1962–2026. Pre-history (grey) · Summer I 1985–1996 · Winter I 1996–2012 · Summer II Oculus 2012–2018 · Trough 2018–2020 · Metaverse summer 2020–2022 · Metaverse winter 2022–2024 · Spatial computing 2024–present.

**VR's cycles run on hardware physics, not algorithms.** The 1985–95 summer collapsed because the silicon could not deliver what the marketing promised — the Virtual Boy is the canonical artefact, but the deeper failure was that PC graphics in 1995 simply could not render an immersive scene at sufficient resolution and framerate without inducing nausea. Sixteen years of Moore's Law had to pass before Oculus could reasonably try again.

**The 2014 Facebook acquisition was a 7-year head fake.** Zuckerberg's \$2 B Oculus bet bought the company a generation of capital, but consumer VR underperformed every analyst forecast between 2016 and 2019. Quest (2019) and Quest 2 (2020) finally cracked the price-and-comfort curve. Then COVID, idle attention, and crypto-fuelled liquidity converged into the metaverse hype peak of late 2021.

**The metaverse winter was largely an AI eclipse.** Cumulative Reality Labs losses crossed \$46 B by 2023. ChatGPT (Nov 2022) didn't just compete for headlines — it competed for engineering talent, capital allocation, and executive attention at every large tech company. Apple's "spatial computing" rebrand in 2024 is the same hardware story under a new label; Vision Pro production was cut back by 2025. *Confidence: high on numbers, moderate on causal weighting of the AI eclipse.*

## Where the cycles correlate, anti-correlate, and where one ate the other.

Two lanes on one axis — read vertically to see what coincided. Period boundaries colour-coded by lane along the bottom.




Figure 3 — AI lane and VR lane on a shared 1943–2026 axis. Bottom rows: AI then VR period boundaries in their season colours.

**AI has a twenty-five-year head start.** McCulloch & Pitts published the first computational neuron the same year ENIAC was being built, while VR's lane is empty until Sutherland's HMD in 1968 and only populates from 1984. AI is fundamentally an older field with a deeper bench of theoretical results.

**Read vertically.** The 1990s are the only true *shared* winter — expert systems collapsed and Virtual Boy flopped within two years. Both fields then disappeared from headlines for ~15 years while underlying tech improved. **2012 was the shared ignition year:** AlexNet and the Oculus Kickstarter are within months of each other, both depending on commodity GPUs.

**The 2022–24 anti-correlation is the most consequential dynamic of the decade.** ChatGPT arriving thirteen months after Facebook's Meta rebrand redirected almost every relevant capital pool — venture, talent, executive attention, GPU allocation — toward LLMs and away from headsets. The metaverse winter is real but only partially endogenous; a large component is exogenous. *Confidence: high on the timing, moderate on the causal magnitude.*

## Three lessons from the chart.

### I · Different physics

AI cycles are funding-and-credibility cycles around algorithmic claims. VR cycles are hardware-and-adoption cycles around display, optics, and silicon. Treating them as the same kind of bubble is an analytical error.

### II · 2012 was not coincidence

Both fields ignited within months of each other because both had been waiting for the same enabling technology: commodity GPUs cheap and powerful enough to use outside specialised contexts. The same firms — Facebook, Google — made parallel bets within two years. When the underlying technology base shifts, multiple fields move in lock-step.

### III · The eclipse

Hype is a near-zero-sum resource. When ChatGPT arrived in late 2022, it didn't merely outshine the metaverse — it cannibalised the capital, GPUs, and engineering attention that the metaverse pitch had just claimed. Spatial computing's 2024 rebrand is in part an attempt to escape that gravity well.

---

## Dr. George Papagiannakis

*Founder & CEO, ORamaVR · Principal Researcher, FORTH-ICS · Professor, University of Crete*

---

*Compiled May 2026 · sourced from primary records where verifiable · confidence flagged inline · prose interprets, charts report*

## All forty-two events plotted on the AI and VR timelines.

The same data that appears as hover-tooltips in the interactive HTML version. AI events first (twenty-six), then VR / metaverse events (sixteen). Each row: lane, year or date, event title, and the explanation.

LANE	YEAR / DATE	EVENT	WHAT IT WAS & WHY IT MATTERED
<b>ARTIFICIAL INTELLIGENCE – 1943 to 2026</b>			
<b>AI</b>	1943	<b>McCulloch &amp; Pitts</b>	First mathematical model of neural computation. Proved that networks of binary threshold neurons can compute any logical function — the seed of every later neural model from the Perceptron through AlexNet to Transformers. Published the same year ENIAC was being built.
<b>AI</b>	1949	<b>Hebb's learning rule</b>	Donald Hebb proposed that synaptic strength increases when neurons fire together. The first credible biological learning rule — directly inspires modern neural network learning algorithms decades later.
<b>AI</b>	1950	<b>Turing — Computing Machinery &amp; Intelligence</b>	Turing proposes the imitation game and argues machines could think in principle. Frames the philosophical question — what does it mean for a machine to be intelligent? — that the field will inherit and never quite resolve.
<b>AI</b>	summer 1956	<b>Dartmouth Workshop</b>	Two-month workshop at Dartmouth College, organized by John McCarthy with Minsky, Shannon, Newell, Simon, and Rochester. McCarthy coins the term 'artificial intelligence' here. Often cited as the field's

LANE	YEAR / DATE	EVENT	WHAT IT WAS & WHY IT MATTERED
			founding event — though the science was already in place; what Dartmouth added was a name, a programme, and a funding pitch.
AI	1957	<b>Rosenblatt — Perceptron</b>	Frank Rosenblatt's pattern recognition algorithm, implemented as the Mark I Perceptron hardware. Generated enormous press attention and military funding. Becomes the canonical example of AI overpromise after Minsky & Papert's 1969 critique.
AI	1958	<b>McCarthy — LISP</b>	John McCarthy creates LISP at MIT. Becomes the dominant AI programming language for 30+ years. Its expressive power for symbolic computation shapes how AI researchers think about the problem.
AI	1965	<b>Simon's 20-year prediction</b>	Herbert Simon famously predicts: "Machines will be capable, within twenty years, of doing any work a man can do." The defining moment of AI overpromise — the prediction expires in 1985 with no breakthrough, contributing to the second AI winter narrative.
AI	1966	<b>Weizenbaum — ELIZA</b>	Joseph Weizenbaum's chatbot. Pattern-matched scripts simulating a Rogerian psychotherapist. Many users believed it was sentient — Weizenbaum himself was alarmed by this and later wrote critical books about the field.

LANE	YEAR / DATE	EVENT	WHAT IT WAS & WHY IT MATTERED
AI	1969	<b>Minsky &amp; Papert — Perceptrons</b>	Book proves that single-layer perceptrons cannot learn the XOR function. Widely (and unfairly) interpreted as proving neural networks have fundamental limits. Effectively kills mainstream neural network research for ~17 years until backpropagation.
AI	1973	<b>Lighthill Report</b>	James Lighthill's report to the UK Science Research Council concludes AI has failed to deliver on its promises. UK government largely withdraws funding from AI research. The first national-scale loss of confidence.
AI	1974	<b>DARPA Speech Understanding cancelled</b>	DARPA terminates the Speech Understanding Research program after disappointing results from CMU and SRI. Marks the end of the post-Sputnik decade of generous US AI funding.
AI	1980	<b>DEC's XCON expert system</b>	Digital Equipment Corporation's expert system for configuring VAX computers. Estimated to save DEC ~\$40M per year by 1986. First convincing proof that expert systems can pay for themselves — triggers the second AI summer and a wave of expert-system startups.
AI	1986	<b>Backpropagation popularised</b>	Rumelhart, Hinton & Williams publish their Nature paper on backpropagation for training multi-layer networks. The algorithm itself was older (Werbos 1974), but this

LANE	YEAR / DATE	EVENT	WHAT IT WAS & WHY IT MATTERED
			paper makes it usable. Solves the XOR problem and quietly reignites neural network research.
AI	1987	<b>LISP machine market collapse</b>	Specialized LISP machine vendors — Symbolics, LMI, TI Explorer — collapse as Sun and Macintosh workstations match their performance at a fraction of the price. The trigger event of the second AI winter.
AI	1992	<b>Japan's Fifth Generation project ends</b>	Japan's \$400M+ Fifth Generation Computer Systems project ends after 10 years with no commercial breakout product. Definitive end of the expert-systems era; the field disperses into machine learning, statistics, and other rebranded forms.
AI	1997	<b>Deep Blue beats Kasparov</b>	IBM's chess computer defeats world champion Garry Kasparov in a six-game match. First time a reigning champion lost to a machine in tournament conditions. Brute-force search and hand-tuned evaluation, not 'intelligence' as the public imagined — but a credibility milestone.
AI	2006	<b>Hinton — Deep Belief Networks</b>	Hinton, Osindero & Teh's deep-belief-network paper rebrands neural networks as 'deep learning.' Demonstrates that layer-wise pretraining solves vanishing-gradient problems. Quiet at the time,

LANE	YEAR / DATE	EVENT	WHAT IT WAS & WHY IT MATTERED
			but the rebrand sticks.
AI	2009	<b>Fei-Fei Li launches ImageNet</b>	14M-image hand-labeled dataset organised around WordNet. Provides the benchmark on which the deep learning revolution will be measured. Without ImageNet, AlexNet doesn't have a stage in 2012.
AI	2011	<b>Watson wins Jeopardy!</b>	IBM's Watson defeats Ken Jennings and Brad Rutter on Jeopardy!. Massive PR moment for AI generally. The underlying tech (DeepQA) is largely classical NLP — but the public read it as machines winning at human knowledge tasks.
AI	2012	<b>AlexNet wins ImageNet</b>	Krizhevsky, Sutskever & Hinton win ImageNet with a deep CNN trained on two consumer GTX 580 GPUs. Halves the error rate of the previous best entry. THE ignition event of the deep-learning era — every subsequent breakthrough builds on this template.
AI	2014	<b>Goodfellow — Generative Adversarial Networks</b>	Two networks compete: one generates samples, one tries to discriminate real from generated. Foundation for modern image and video generation; predecessor to today's diffusion models in the public imagination.
AI	2016	<b>DeepMind — AlphaGo beats Lee Sedol</b>	DeepMind's AlphaGo defeats Lee Sedol 4–1. Go was long thought to be a decade away. Game 4 was a famous

LANE	YEAR / DATE	EVENT	WHAT IT WAS & WHY IT MATTERED
			human win; the rest were AlphaGo dominance — including Move 37 in Game 2, considered superhuman.
AI	2017	<b>Vaswani et al. — Attention Is All You Need</b>	Self-attention replaces RNN and CNN architectures for sequence modelling. Foundation of every subsequent large language model — BERT, GPT-1/2/3/4, Claude, Gemini. Arguably the most consequential ML paper of the decade.
AI	2020	<b>OpenAI — GPT-3</b>	175-billion-parameter language model. First convincing demonstration of in-context / few-shot learning at scale. Spawns the LLM commercial era; copilot-style applications begin appearing throughout 2021–22.
AI	30 Nov 2022	<b>ChatGPT released</b>	OpenAI releases ChatGPT publicly. Reaches 100M monthly active users within 2 months — the fastest consumer product adoption in history. Generative AI becomes mainstream overnight; cannibalises capital, GPU allocation, and engineering attention from the metaverse.
AI	2024	<b>OpenAI o1 — reasoning at inference</b>	First production model to use inference-time chain-of-thought reasoning at scale. Paradigm shift from pretrained generation toward deliberative reasoning. Followed by Anthropic's extended thinking, DeepSeek-R1, and similar reasoning

LANE	YEAR / DATE	EVENT	WHAT IT WAS & WHY IT MATTERED
			systems.
<b>VIRTUAL REALITY / METAVERSE – 1962 to 2026</b>			
VR	1962	<b>Heilig — Sensorama</b>	Morton Heilig's mechanical immersive cabinet. The first multi-sensory machine: stereo 3D film, stereo sound, smell, vibration, wind. Patented but commercial failure — way ahead of the technology base needed to support it.
VR	1968	<b>Sutherland — Sword of Damocles</b>	Ivan Sutherland's tethered head-mounted display at the University of Utah. First HMD with head-tracking and stereoscopic 3D. So heavy it had to be suspended from the ceiling — hence the name. Conceptual foundation of every later HMD.
VR	1984	<b>Lanier — VPL Research founded</b>	Jaron Lanier founds VPL Research. First commercial VR company. Lanier coins 'virtual reality' as a marketing term to bundle the EyePhone, DataGlove, and DataSuit into a sellable concept.
VR	1989	<b>VPL — EyePhone &amp; DataGlove ship</b>	VPL's commercial VR products: EyePhone HMD (\$9,400–\$49,000) and DataGlove (\$8,800). Hopelessly expensive for any market beyond research labs and high-end industrial — but seeded a generation of VR researchers and shaped the 1990s aesthetic.
VR	1992	<b>Stephenson — Snow Crash</b>	Neal Stephenson's novel introduces 'metaverse' to the language — a

LANE	YEAR / DATE	EVENT	WHAT IT WAS & WHY IT MATTERED
			shared virtual space accessed via avatars. Released the same year as the film <i>The Lawnmower Man</i> . Cultural peak of the first VR hype cycle. Every subsequent metaverse pitch is, consciously or not, descended from this book.
VR	1995	<b>Nintendo — Virtual Boy flops</b>	Nintendo's stereoscopic VR console. Red-on-black monochrome display, induced headaches and nausea, no head-tracking. Commercial flop; discontinued in 1996. Became the canonical artefact of why VR couldn't deliver in the 1990s — the silicon and optics simply weren't ready.
VR	1999	<b>Sun acquires VPL patents</b>	Sun Microsystems acquires VPL Research's patent portfolio after VPL's bankruptcy. Symbolic end of the first commercial VR era. The patents largely sit dormant for a decade.
VR	2003	<b>Linden Lab — Second Life</b>	Linden Lab's screen-based virtual world (no HMD). Reaches ~1M active users by 2007. Not VR proper, but the metaverse-precursor concept proven at consumer scale — and the ancestor of every later social-virtual-world pitch.
VR	Aug 2012	<b>Luckey — Oculus Rift Kickstarter</b>	Palmer Luckey's Kickstarter campaign raises \$2.4M (10× goal). Built on commodity smartphone displays and IMUs — the same hardware wave powering AlexNet a few months later. Reignites VR after

LANE	YEAR / DATE	EVENT	WHAT IT WAS & WHY IT MATTERED
			a 15-year hiatus. THE ignition event for the second VR era.
VR	Mar 2014	<b>Facebook acquires Oculus</b>	Facebook acquires Oculus VR for \$2 billion. Zuckerberg's seven-year bet on VR as the next computing platform. Foreshadows the 2021 Meta rebrand — and same year Google acquires DeepMind, parallel bets across both fields.
VR	2016	<b>Consumer VR ships — Rift, Vive, PSVR</b>	Three consumer VR headsets ship within months: Oculus Rift (\$599), HTC Vive (\$799), PlayStation VR (\$399). First true consumer VR generation. Sales underperform every analyst forecast — the killer app problem becomes visible.
VR	Oct 2020	<b>Meta — Quest 2 launches</b>	Quest 2 ships at \$299 — a price point that finally cracked the consumer adoption curve. Sells ~20M+ units, becoming the best-selling VR headset of all time. COVID lockdowns and idle attention fuel adoption.
VR	28 Oct 2021	<b>Facebook → Meta rebrand</b>	Facebook renames to Meta. Zuckerberg announces the metaverse as the next computing platform and pledges tens of billions in investment. The defining moment of metaverse hype — and, with hindsight, the peak.
VR	2022	<b>Reality Labs loses \$13.7B</b>	Meta's Reality Labs division loses \$13.7 billion in fiscal 2022 alone. Cumulative losses since 2019 exceed

LANE	YEAR / DATE	EVENT	WHAT IT WAS & WHY IT MATTERED
			\$36B by year-end. Public scepticism mounts; Meta share price falls; ChatGPT launches one month later and absorbs the remaining oxygen.
VR	Feb 2024	<b>Apple Vision Pro ships</b>	Apple Vision Pro ships at \$3,499. Rebranded as 'spatial computing' rather than VR/metaverse. Premium positioning, hand+eye tracking, dual 4K micro-OLED panels. Production targets reportedly cut within 12 months as the device searches for its killer app.
VR	2026	<b>Today (May 2026)</b>	Apple has cut Vision Pro production. Meta has pivoted to Quest 3S at \$299 — the low-end mainstream play. The metaverse winter is real but partial; the underlying technology continues to improve while public enthusiasm has migrated to AI.