PATTI INSIGHT LEGENDS · No. 04

If Patti Had Existed Back Then (No. 04): Why Did the GPU Win AI?

Historical retrospective edition

Edition	Verification ID	Data cutoff	Issue date	Publisher
Historical retrospective	Promotional article; no verification ID issued	Act 3: end of 2006 Other sections: queried 2026-11-03	November 3, 2026	Farmar Kaohsiung IP Office

A technology-history retrospective and promotional article. Act 3 is a simulation rebuilt from information public by the end of 2006, not an actual document from the time. Not investment advice; not legal advice.

If Patti Had Existed Back Then (No. 04): Why Did the GPU Win AI?

> Patti Insight, Legends series. We take a piece of technology history everyone thinks they know and ask: what could be seen at the time, what should have been asked, and which signposts were worth watching (signals that a road is starting to open). Every section says plainly what was **visible at the time** and what is **hindsight**.
>
> Codes such as 〔S05〕 correspond to the "Source codes" appendix at the end. The verification status of the main factual sentences is kept in a separate record, available from the publisher on request.

On December 30, 2014, the U.S. International Trade Commission (ITC) opened an investigation into a batch of graphics processing chips (GPUs), systems on a chip (SoCs) and products containing them. The respondents named in the notice were not just NVIDIA but more than a dozen other entities, among them Biostar, Elitegroup, Sparkle, EVGA and ZOTAC. The complainant was Samsung 〔S10〕.

An administrative law judge later found a violation of three patents and recommended a limited exclusion order (barring specific products from import into the U.S.); during the Commission's review the parties settled, so there was no final decision 〔S10〕.

The most interesting thing is the titles of those three patents: CMOS static random access memory, a shared bus system, a data strobe buffer. **In a case about GPUs, judging by the titles, none is about graphics 〔S72〕.**

Just over two months earlier (October 2014), the ITC had opened an investigation of Samsung and Qualcomm on NVIDIA's complaint; the asserted patents included a graphics processor patent, US 7,038,685 (the '685). NVIDIA had asserted seven and later withdrew four, and the rest ended in a finding of no violation 〔S02〕〔S09〕. Whether one case caused the other, the documents don't say.

You have read the GPU story many times: near-death, the bet on CUDA (NVIDIA's GPU programming platform), the turn at AlexNet (the neural network that won an image-recognition contest in 2012). Others tell it superbly, and we won't retell it. We ask one thing the general histories rarely dwell on: **what did patents actually do?**

We found three things.

First, in the cases we read, patents affected who paid whom, and who had no license. In 2011 Intel agreed to pay NVIDIA $1.5 billion in licensing fees; NVIDIA's license did not cover certain microprocessors (defined in the agreement as Intel Processors or Intel Compatible Processors), certain chipsets that connect to Intel Processors, or certain flash memory products; the same agreement also settled litigation, so not all of it can be credited to patents 〔S05〕. In 2010 NVIDIA took a limited exclusion order in the Rambus case and also signed a license: 1% of the net sales price per unit for certain memory controllers and 2% for certain others 〔S06〕.

Second, the patents NVIDIA ran into were not necessarily graphics-architecture patents. The Rambus case concerned memory controllers, and Samsung's three aren't about graphics either, judging by titles (we did not read each claim, the part that states what a patent actually covers). Among the cases we read, the two in which an exclusion order issued or was recommended were not about graphics architecture (our reading) 〔S06〕〔S72〕.

Third, no document we read says the GPU won AI through patents. Research papers don't usually discuss patents, so this means we can't tell, not that it is disproved 〔S18〕〔S19〕〔S20〕〔S21〕.

Act 1 — Take the patents away: would it change?

This table is the core of the article. It asks the question event by event.

Event	What the documents say	Without patents, would it change? (our judgment)
Rambus (2008–10)	The ITC found that NVIDIA infringed three patents and issued a limited exclusion order naming NVIDIA and several graphics-card-related companies; NVIDIA signed a license at 1%/2% royalties 〔S06〕〔S73〕〔S74〕	**Yes:** cost and import risk. This is the one case in these documents where an exclusion order actually issued; we did not check what followed
Intel (2009–11)	Intel sued first, over whether a 2004 license covered certain chipsets; cross license, Intel pays $1.5 billion 〔S05〕〔S06〕	**Yes:** cash and boundaries both come from the license. Why NVIDIA stopped developing future chipsets isn't stated in what we read, so we can't say it was forced out 〔S08〕
NVIDIA sues Samsung and Qualcomm (2014–15)	No violation; two claims of the '685 not held invalid, but not infringed 〔S09〕	**Not for the rivals:** these patents didn't stop the accused products
Samsung sues NVIDIA and others (2014–16)	Judge found violation, recommended an exclusion order; settled during review 〔S10〕	**Yes:** but it stopped at a recommendation, with no final decision
Research papers (2010–13)	Full-text search of four papers: "patent" and "license" appear zero times 〔S18〕〔S19〕〔S20〕〔S21〕	**Can't tell:** papers don't usually discuss them
CUDA	The visible protection is a license term: no reverse engineering output in order to translate it to non-NVIDIA platforms (current terms; the terms of the time were not checked) 〔S17〕	**Can't tell from what we read:** this is contract, not patent

Settlements and private licenses are invisible, so this table can only ask about the public part.

The 15 claims we drew up. These are possible claims we thought of ourselves, not all necessarily made by anyone; the full table with each check is available from the publisher on request. "NVIDIA crushed 3dfx with patents": the documents show 3dfx sued first in 1998, NVIDIA countersued on five patents in 2000, later bought the assets, patents included, for about $70 million, and both dropped their suits; we read nothing tying 3dfx's decline to patents 〔S07〕. "Intel paid $1.5 billion for infringing": those were licensing fees under a cross license, paid in six installments 〔S05〕.

“Patents are NVIDIA’s moat”: not a word from the documents, and NVIDIA’s annual report carries boilerplate risk language that patents “may not provide us with competitive advantages” 〔S06〕.



The 15 claims we drew up: how they checked out

Figure: Of the 15, 5 were wrong, 5 partly right or needing correction, 2 had no primary source, and 3 were unverifiable attributions; items were chosen by the authors, so this result is not a statistic. Sources for the three examples: 〔S05〕〔S06〕〔S09〕〔S19〕.

Act 2 — How someone back then would have asked

Set the clock to the end of 2006. You run engineering at a mid-sized board maker. NVIDIA has just announced CUDA, and people are also talking about multicore and special-purpose accelerators. Your CEO asks: should we get into high-performance computing?

"What are the trends in GPU computing?" will mostly get you an encyclopedia-style overview. Patti would change it to: **We are a mid-sized board maker deciding whether to commit to compute boards over the next five years; where is each of three routes stuck; what signals would change my decision; who holds the patents on this road, and could I run into them; use only information public by the end of 2006.**

(Illustration. In practice Patti first restates and reorganizes your question; the restatement screen keeps your original question word for word, then lists suggested aspects for you to tick.)

Act 3 — The 2006 report (data cutoff: end of 2006)

The following was rebuilt by hand by the authors along these lines; it is not output from a real Patti run. It uses only information public by the end of 2006; one survey published slightly later, describing the state of things then, is marked. The data keeps to the cutoff, but which documents we picked and which signposts we set are shaped by hindsight. A patent's "issue" date is when the office made it public.

The problem. In December 2006, a Berkeley report wrote of the power wall and the memory wall (limits set by chip power and memory speed), and quoted Sterling's question at the SC06 conference: "will multicore ultimately be asphyxiated by the memory wall?" 〔S24〕. Where would computing power come from?

Route 1: Multicore CPUs. Multicore CPUs were the mainstream of the time, stuck on power, memory and the parallel programming model 〔S24〕.

- Signpost (set by the authors): Ordinary software gets multicore speedups without major rewriting; five years.

Route 2: General-purpose GPU computing. Ian Buck's 2006 thesis used Brook (a programming language that extends C) to treat the GPU as a streaming co-processor 〔S37〕. That November, NVIDIA announced CUDA, a C-compiler environment for the GPU, claiming speeds up to 100 times faster than traditional approaches (a company claim) 〔S27〕. Of the hardware before it, a survey said "The GPU is hardly a

computational panacea": it lacked double precision and integers, and problems had to be recast in graphics terms (published in 2007, describing earlier hardware) 〔S23〕.

- Signpost (set by the authors): An independent report of a non-graphics application 10 times faster or more, without recasting as graphics; five years.

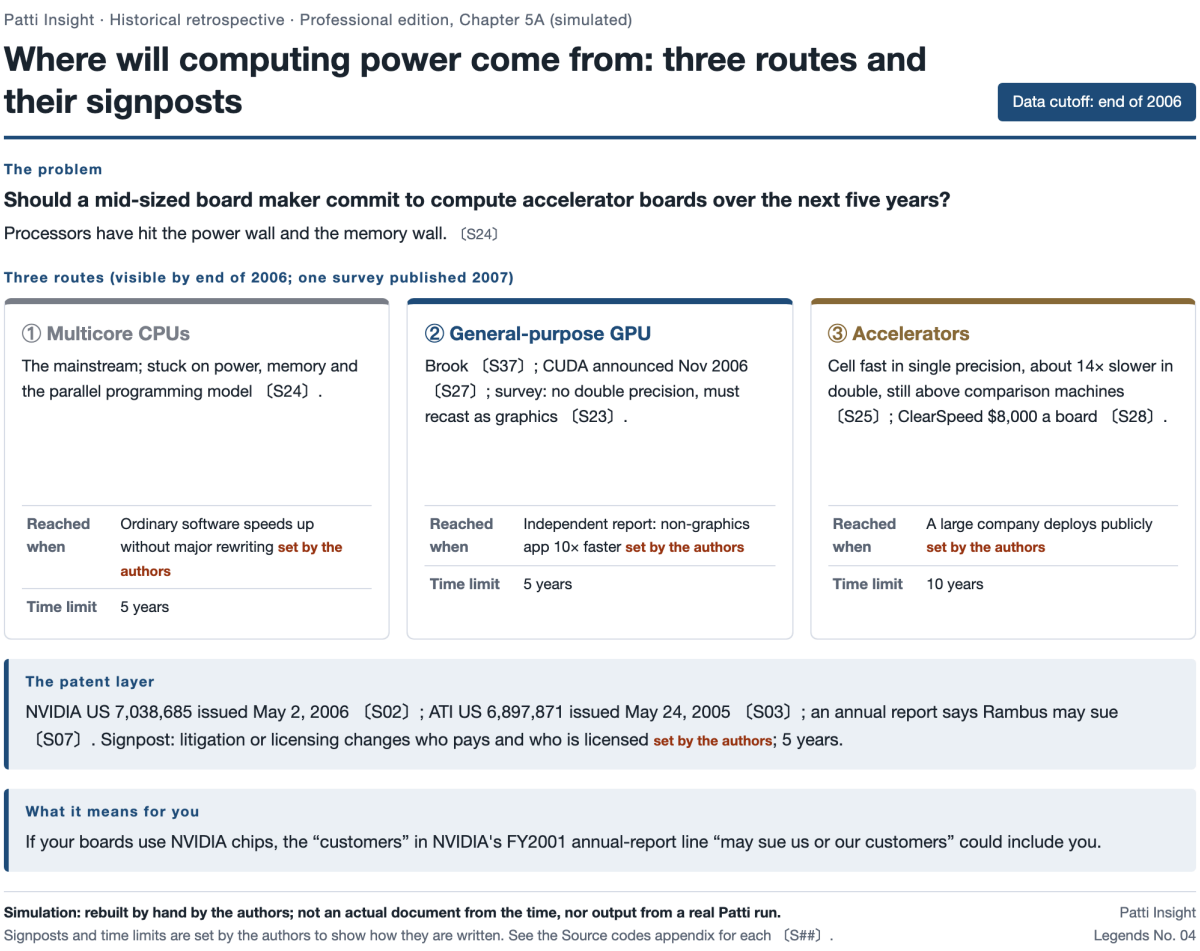
Route 3: Special-purpose accelerators. The Cell processor (designed by a partnership of Sony, Toshiba and IBM) was extremely fast in single precision, but researchers called single precision "widely considered insufficient for the majority of scientific applications," and double precision took about 14 times as long, though it still beat the comparison machines of the day 〔S25〕. ClearSpeed cost $8,000 per board; a press report said its products must run at least 5 times faster than the best general-purpose chips from Intel and AMD 〔S28〕.

- Signpost (set by the authors): A large company publicly deploying a special-purpose accelerator; ten years.

The patent layer. NVIDIA's US 7,038,685, issued May 2, 2006, covers one processing unit handling several sample types with multiple threads, balancing work cycle by cycle; ATI's US 6,897,871, issued May 24, 2005, has "unified shader" in its title 〔S02〕〔S03〕. Both are graphics patents about one unit handling different work. Our reading is that this is close to general-purpose computing; neither first claim mentions general-purpose computing. We make no call on priority or infringement. NVIDIA's FY2001 annual report had warned that Rambus "may sue us or our customers" (we did not check whether the 2006 report still said so) 〔S07〕.

- Signpost (set by the authors): Litigation or licensing among the major players that changes who pays and who is licensed; five years.

What it means for you. If your boards use NVIDIA chips, the "customers" in NVIDIA's annual-report line "may sue us or our customers" could include you.



The one-page version of Act 3 (simulated; data cutoff end of 2006)

Figure: The one-page version of Act 3, rebuilt by hand by the authors 〔S02〕〔S03〕〔S07〕〔S23〕〔S24〕〔S25〕〔S27〕〔S28〕〔S37〕. It is a simulation, not an actual document from the time; signposts and time limits are our own.

> **If Patti had existed back then: is this table enough to decide?** Would seeing these four signposts change your decision? **What does it miss?** Think it over, then read on.

Act 4 — The reveal (hindsight)

Everything in this section was learned after the fact.

What the Act 3 table missed was in the documents it cited. That Berkeley report calls machine learning “One of the most promising areas,” but its examples are support vector machines and principal component analysis, filed under linear algebra; a full-text search finds the word “neural” zero times 〔S24〕. Yet the Owens et al. survey that Act 3 cites mentions someone computing neural networks on graphics hardware in 1998 〔S23〕 (we read the version published in 2007; we did not check whether the 2005 version has this sentence); and according to a 2010 paper, people did it in 2005 as well, with modest speedups then 〔S19〕. Bryan Catanzaro was among the Berkeley report’s authors. Seven years later his name appears on another paper, under NVIDIA: a billion parameters trained on three GPU machines, set against an earlier result that used 16,000 CPU cores 〔S21〕.

The literature of the time named double precision as a bottleneck, but the 2010 handwriting-recognition application didn’t need it. In 2007 AMD announced a double-precision GPU (its own account) 〔S32〕. But what broke the handwriting-recognition record in 2010 was “single precision floating-point GPU-based neural nets,” which trained more than 40 times faster than on standard microprocessors 〔S19〕.

Google discussed GPUs in 2006 and passed. It later wrote that it had discussed GPUs, FPGAs and ASICs, but idle datacenter capacity sufficed: “it’s hard to improve on free” 〔S22〕. At the 2012 NIPS conference (the annual machine-learning meeting) there were two roads: AlexNet trained for five to six days on two GTX 580s, its paper said “Luckily, current GPUs … are powerful enough,” and it won the ILSVRC-2012 image-recognition contest with a top-5 error rate (the share of images whose correct label is not among the model’s five top guesses) of 15.3% against 26.2% for the second-best entry 〔S18〕; Google used a CPU cluster of tens of thousands of cores and said the speed-up was small when a model didn’t fit in GPU memory 〔S20〕. So the GPU didn’t win on the spot. Google’s conversation changed only in 2013, because of a sum: if people used voice search for three minutes a day, its datacenters would have to double. The result was its own ASIC for inference, with off-the-shelf GPUs bought for training 〔S22〕.

CUDA wasn't born for AI either. The applications named at its 2006 launch were product design, data analysis, technical computing and game physics 〔S27〕. NVIDIA's FY2015 annual report names deep learning explicitly (we did not compare earlier reports in full) 〔S15〕.

How did the table do?

Signpost (set by the authors)	Time limit	What we read	Result
① Ordinary software gets multicore speedups	2011	We did not check this outcome	Not checked
② Independent report: non-graphics application 10 times faster	2011	By 2010 at the latest, using CUDA, Cireşan and colleagues trained neural networks over 40 times faster than on standard microprocessors 〔S19〕	Reached
③ A large company publicly deploys a special-purpose accelerator	2016	Microsoft's Catapult published in 2014, deployed in 2015; Google's TPU deployed in 2015 〔S22〕	Reached
④ Litigation or licensing changes who pays and who is licensed	2011	Rambus 2010, Intel 2011 〔S05〕〔S06〕	Reached

We set these signposts after the fact; this only shows how such a table is written and proves no decision right.

Back at your board maker. Signposts ② and ④ didn't appear until 2010–11, just as your five-year decision came due. But that "our customers" didn't wait. In December 2008 the ITC opened an investigation on Rambus's complaint; the notice says the products include "graphics cards and motherboards," and besides NVIDIA it names ASUS, BFG, Biostar, EVGA, Giga-byte, MSI, Palit, Sparkle and others 〔S73〕. The limited exclusion order and cease-and-desist orders of July 2010 name ASUS, BFG, Biostar, EVGA, MSI and Palit, among others 〔S74〕. On Samsung's 2014 list, Biostar, EVGA and Sparkle appeared again 〔S10〕. As a board maker, you could receive such

a notice too (“board maker” is our shorthand; the notice does not call each of them that).

Not on the table: export controls. In August 2022 the U.S. government required licenses to export A100 and H100 to China and Russia 〔S12〕; the October rule set thresholds such as 600 Gbyte/s 〔S11〕. This thread has nothing to do with patents, yet it decides directly who can be sold to.

Where the documents don’t mention patents. Google built its TPU, and its paper even points readers to its own patent applications 〔S22〕; when Intel canceled the consumer version of Larrabee (its own graphics chip), it said “silicon and software development are behind” 〔S31〕. We read nothing showing patents played a part in either, which doesn’t rule it out.

What we could not check. We did not read the full text of the 2009 paper by Raina et al. on deep learning with graphics processors, Brook or Buck’s thesis, the ImageNet paper, the FPGA route, documents on NVIDIA’s internal pivot (we saw only book-review excerpts), the public ITC opinions, the claims of each Samsung and Rambus patent, or the book on NVIDIA by Stephen Witt; nor did we check the multicore signpost’s outcome.

For engineers today: in the cases we read, the tolls are real, and the toll booth may not be on the layer you are watching.

Interlude — When history can reverse (hindsight)

Everything in this section is hindsight.

By 2017, challengers were already on the field. Google built its TPU for inference, aiming for 10 times the cost-performance of GPUs, while still buying GPUs for training (its own 2017 assessment, which can’t be extended to today) 〔S22〕.

Patents change hands too. In November 2024, roughly 100 U.S. patents moved from AMD via ATI to Onesta IP, which then sued NVIDIA; we did not check the outcome

〔S36〕. Someone else's patent can come calling under a new owner.

> **If you have Patti today:** in your class, who is filing, and who isn't? The next act shows an example.

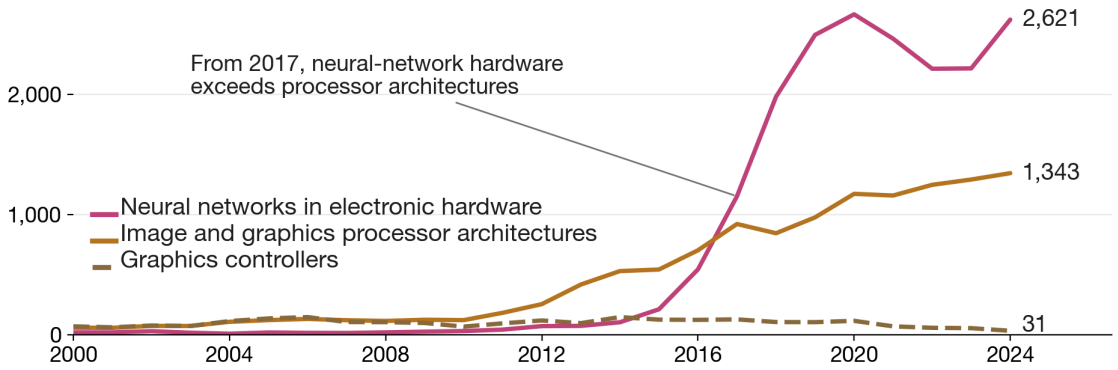
Act 5 — Who is filing?

Patent offices tag each patent with technical classes. We count, year by year, one invention's filings across countries (a patent family), then look at who the applicant is.

Patent activity from graphics to AI
By classification, not citation

New patent families per year. The subject is patent activity, not market size or who won.

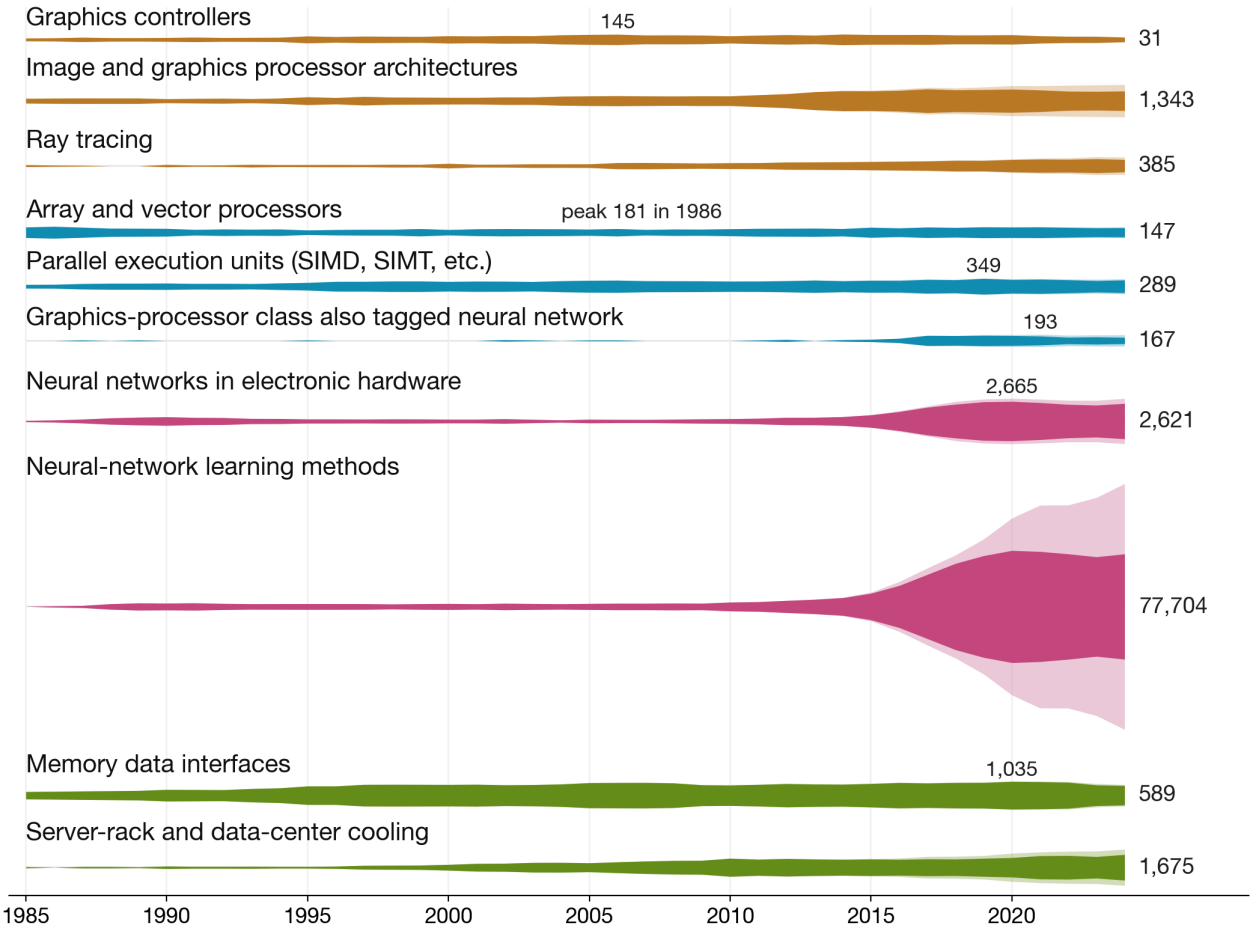
Graphics processors and neural-network hardware, families per year



By classification, 1985–2024 (band width = square root of families; one scale for all bands)

Graphics | Parallel and overlap | Neural networks | Memory and cooling

Dark: families filed outside China too. Light: families filed only in China.



Data: Google Patents Public Data (BigQuery), queried 2026-11-03; latest publication date 2026-09-24. Families use earliest priority year. 2024 may still be incomplete; 2025 onward is not drawn. Classes follow CPC; older documents were mostly classified retroactively. A family can fall in several classes, so they do not add up. All patent families also rose in 2000–2024, mostly filed in China; width does not mean technical

importance. Classes are not the same as GPUs: processor architectures include general image processors. Stacked packaging is not drawn: data coverage is inadequate. Chart: Patti Insight.

Patent activity from graphics to AI: by classification, not citation

Figure: New patent families per year, by each family's earliest priority year〔S70〕. Branches are assigned by patent classification — **not by citation**. The subject is **patent activity**, not the market, and not who won. A family can belong to several branches, so branches can't be added up. 2024 may not be fully published; all patents grew in this period too, largely in China. Sampled titles show that many in the neural-network hardware branch don't reveal any hardware content, and a class isn't the same as a GPU.

2017 is the crossover. From that year, families in neural-network hardware outnumber image and graphics-processor architecture, even counting only outside China (987 against 716)〔S70〕. But these classes were often added later, and the neural-network subclasses are a recent refinement, so treat the year as approximate.

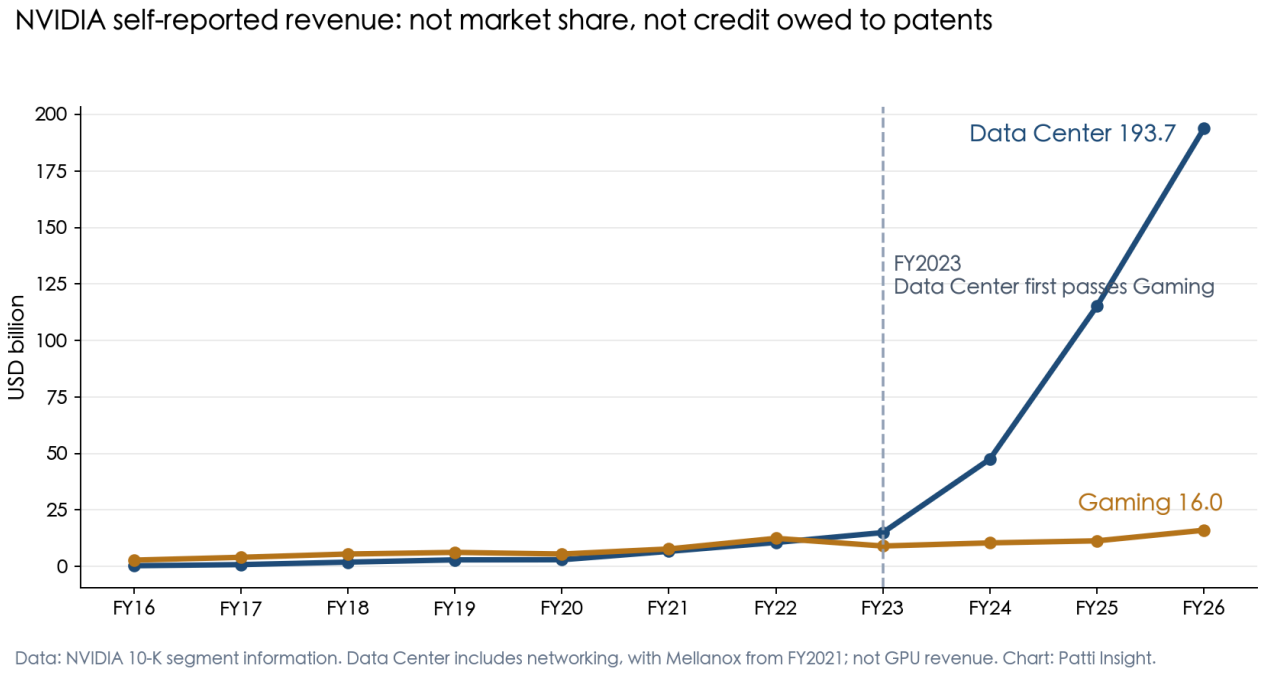
Applicants.

Class	Period	NVIDIA	Other applicants, examples (Samsung at group level; the two Sony companies not combined)
Image and graphics processor architecture	2000–04	55 families, 15.2%	Two Sony companies, 12 and 11 (Sony Computer Entertainment, Sony); Intel 12, AMD/ATI 11
Same	2005–09	87 families, 14.4%	Samsung 42, Intel 29
Same	2010–14	105 families, 7.0%	Intel 161, Samsung 124
Same	2015–19	94 families, 2.4%	Intel 529, Samsung 208
Same	2020–24	188 families, 3.0%	Intel 241, Huawei 172
Neural-network hardware	2020–24	248 families, 2.0%	Samsung 708, Huawei 393, Intel 323

In 2000–09, NVIDIA filed the most in the image and graphics processor architecture class; the class isn't the same as GPUs and also includes image processors. Most of the fall in its share came in 2010–19 and had little to do with China; only in 2020–24 is the gap clearly affected by families filed only in China; excluding those families, NVIDIA's share is about 6.9%, and about 3% in neural-network hardware for the same period〔S70〕. This data supports neither "won on patent count" nor "lost on patent

count"; it covers only some classes, and counts aren't value. In 2015–19 Intel had 529 families and NVIDIA 94: that is filing volume, not cause.

The market layer. NVIDIA's annual reports show Data Center revenue of $339 million in FY2016 and $193.7 billion in FY2026; it first passed Gaming in FY2023 ($15.0 billion against $9.07 billion); Gaming also grew, from $2.82 billion to $16.0 billion 〔S71〕. This includes networking ($31.4 billion in FY2026), and from FY2021 it is also affected by Mellanox (networking equipment), acquired in 2020, so it isn't GPU revenue, and it isn't market share.



NVIDIA Data Center and Gaming revenue, FY2016–FY2026

Figure: Self-reported revenue from NVIDIA's 10-K; fiscal 2026 falls mostly in 2025. Not market share, and it can't tell you who deserves credit 〔S71〕.

Over the same period, NVIDIA's families in the neural-network hardware class also rose, but at only about a 2% share, while Data Center made up nearly nine-tenths of FY2026 revenue. These two numbers can't explain each other: revenue can't be credited to patents, and patents aren't decoration either, as the table in Act 1 shows. Patti answers "who is filing, and where?" — not "who gets the credit for this money?"

Closing: if you have a technical question of your own

The questions in this article can be reused: what could be seen at the time, what should have been asked, and which signposts were worth watching.

If you want Patti Insight to look at your own topic, there are two routes. Ready-made reports are on topics we chose, and you can buy them directly. A tailored study starts with an application: a person reviews your identity and topic first, we send a payment link only if the application is accepted, and research starts after payment. A topic may be declined, and a report may not answer everything you wanted to know. Where Patti hasn't read the full text, it won't judge whether a technical gap (a place no one has solved yet) exists.

This article did not investigate the class you work in.

How Patti Insight works, and its limits: https://patti.empryon.ai/methods

---

Act 3 of this article is a simulation: rebuilt by hand by the authors from public information, not an actual document from the time and not output from a real Patti run. All signposts were set by the authors. Not investment advice; this article does not discuss share prices or market value.

Appendix: Source codes

Each code 〔S##〕 in the text corresponds to a bibliographic line below. Every code used in this article is listed. Entries marked “Not checked” are sources we could not read.

〔S02〕 John Erik Lindholm | “Programmable graphics processor for multithreaded execution of programs,” US 7,038,685 | https://patents.google.com/patent/US7038685B1/en | Primary source (patent)

〔S03〕 Morein et al. (ATI) | “Graphics processing architecture employing a unified shader,” US 6,897,871 | https://patents.google.com/patent/US6897871B1/en | Primary source (patent)

〔S05〕 NVIDIA Form 8-K (January 10, 2011) | https://www.sec.gov/Archives/edgar/data/1045810/000119312511005134/d8k.htm | Primary source

〔S06〕 NVIDIA Form 10-K, FY2011 | https://www.sec.gov/Archives/edgar/data/1045810/000104581011000015/fy2011form10k.htm | Primary source

〔S07〕 NVIDIA Form 10-K405, FY2001 | https://www.sec.gov/Archives/edgar/data/0001045810/000101287001500492/d10k405.txt | Primary source

〔S08〕 NVIDIA Form 10-K, FY2012 | https://www.sec.gov/Archives/edgar/data/1045810/000104581012000013/nvda-2012×10k.htm | Primary source

〔S09〕 USITC | Inv. No. 337-TA-932, Federal Register 2015-31816 | Primary source

〔S10〕 USITC | Inv. No. 337-TA-941, Federal Register 2016-14657 | Primary source

〔S11〕 BIS | Implementation of Additional Export Controls, 87 Fed. Reg. 62186 (October 13, 2022) | Primary source

〔S12〕 NVIDIA Form 8-K (August 26, 2022) | https://www.sec.gov/Archives/edgar/data/1045810/000104581022000146/nvda-20220826.htm | Primary source

〔S15〕 NVIDIA Form 10-K, FY2015 | https://www.sec.gov/Archives/edgar/data/1045810/000104581015000036/nvda-2015×10k.htm | Primary source

〔S17〕 NVIDIA CUDA EULA | https://docs.nvidia.com/cuda/eula/index.html | Primary source

〔S18〕 Krizhevsky, Sutskever, Hinton | ImageNet Classification with Deep Convolutional Neural Networks | NIPS 2012 | Primary source

〔S19〕 Cireşan et al. | arXiv:1003.0358 (2010) | https://arxiv.org/pdf/1003.0358 | Primary source

〔S20〕 Dean et al. | Large Scale Distributed Deep Networks | NIPS 2012 | Primary source

〔S21〕 Coates et al. | Deep learning with COTS HPC systems | ICML 2013 | https://proceedings.mlr.press/v28/coates13.pdf | Primary source

〔S22〕 Jouppi et al. (Google) | In-Datacenter Performance Analysis of a Tensor Processing Unit | arXiv:1704.04760 | Primary source

〔S23〕 Owens et al. | A Survey of General-Purpose Computation on Graphics Hardware | Computer Graphics Forum, 2007 | Scholarly secondary source

〔S24〕 Asanovic et al. | The Landscape of Parallel Computing Research: A View from Berkeley | UCB/EECS-2006-183 | Scholarly secondary source

〔S25〕 Williams et al. | The Potential of the Cell Processor for Scientific Computing | Scholarly secondary source

〔S27〕 Digital Engineering | NVIDIA Unveils CUDA | November 2006 | https://www.digitalengineering247.com/article/nvidia-unveils-cuda/engineering-computing | General media (company release)

〔S28〕 The Register | ClearSpeed | November 18, 2006 | https://www.theregister.com/2006/11/18/clearspeed_silicon_sc06/ | General media

〔S31〕 GamesBeat | Intel cancels Larrabee consumer graphics chip | https://gamesbeat.com/intel-cancels-larrabee-consumer-graphics-chip/ | General media

〔S32〕 PC Perspective, reposting an AMD release | November 8, 2007 | Company account

〔S36〕 Legal500 | Onesta IP | General media

〔S37〕 hgpu.org | Ian Buck 2006 Stanford thesis page | https://hgpu.org/?p=5611 | General media (aggregator)

〔S70〕 Google Patents Public Data (BigQuery public dataset `patents-public-data.patents.publications`) | Queried November 3, 2026 | Primary source (patent bibliographic dataset)

〔S71〕 NVIDIA Form 10-K segment information, FY2016–FY2026 | Primary source (self-reported)

〔S72〕 Google Patents | US 6,147,385 (CMOS static random access memory devices), US 6,173,349 (Shared bus system with transaction and destination ID), US 7,804,734 (Data strobe buffer and memory system including the same), Samsung | Read November 3, 2026 | Primary source (patent bibliography)

〔S73〕 USITC, Inv. No. 337-TA-661, notice of investigation, Federal Register E8-29162 (December 10, 2008, 73 FR 75131) | https://www.govinfo.gov/content/pkg/FR-2008-12-10/pdf/E8-29162.pdf | Primary source

〔S74〕 USITC, Inv. No. 337-TA-661, final determination notice, Federal Register 2010-18715 (July 30, 2010) | https://www.govinfo.gov/content/pkg/FR-2010-07-30/html/2010-18715.htm | Primary source

EMPRYON

Patti Insight

Mechanism-driven technology gap and market reports

Publisher: Farmar Kaohsiung IP Office

Contact data@empryon.ai

Legends No. 04 · Historical retrospective · Promotional article; no verification ID issued · November 2, 2026

Patti Insight is a product of the Empryon brand; this article is published by Farmar Kaohsiung IP Office. Technology-history retrospective; not investment advice, not legal advice. © 2026 Farmar Kaohsiung IP Office. All rights reserved.