



Who is NVIDIA? Click [here](#)→

The New Data Centers Era

Javier Pemán Pérez-Serrano - Director comercial NVIDIA Ltd
Supercomputacion, Centros de investigacion y universidades en Iberia

OCT 2025

Javier Pemán Pérez-Serrano

Ingeniero Industrial- Tecnun San Sebastian

MBA IESE

netConnections (UK) – ASI (USA) – PwC (Madrid) –

IBM (BCN) – NVIDIA (BCN)

Director commercial NVIDIA Ltd Supercomputacion, Centros de
investigacion y universidades en Iberia, Saudi

javierp@nvidia.com

+34 635520529



EL MOMENTO IPHONE DE LA IA

AI Evolution



PHYSICAL AI

AUTONOMOUS VEHICLES
GENERAL ROBOTICS



AGENTIC AI

CODING ASSISTANT
CUSTOMER SERVICE
PATIENT CARE

GENERATIVE AI

DIGITAL MARKETING
CONTENT CREATION

PERCEPTION AI

SPEECH RECOGNITION
DEEP RECSYS
MEDICAL IMAGING

"FIRST CONTACT"

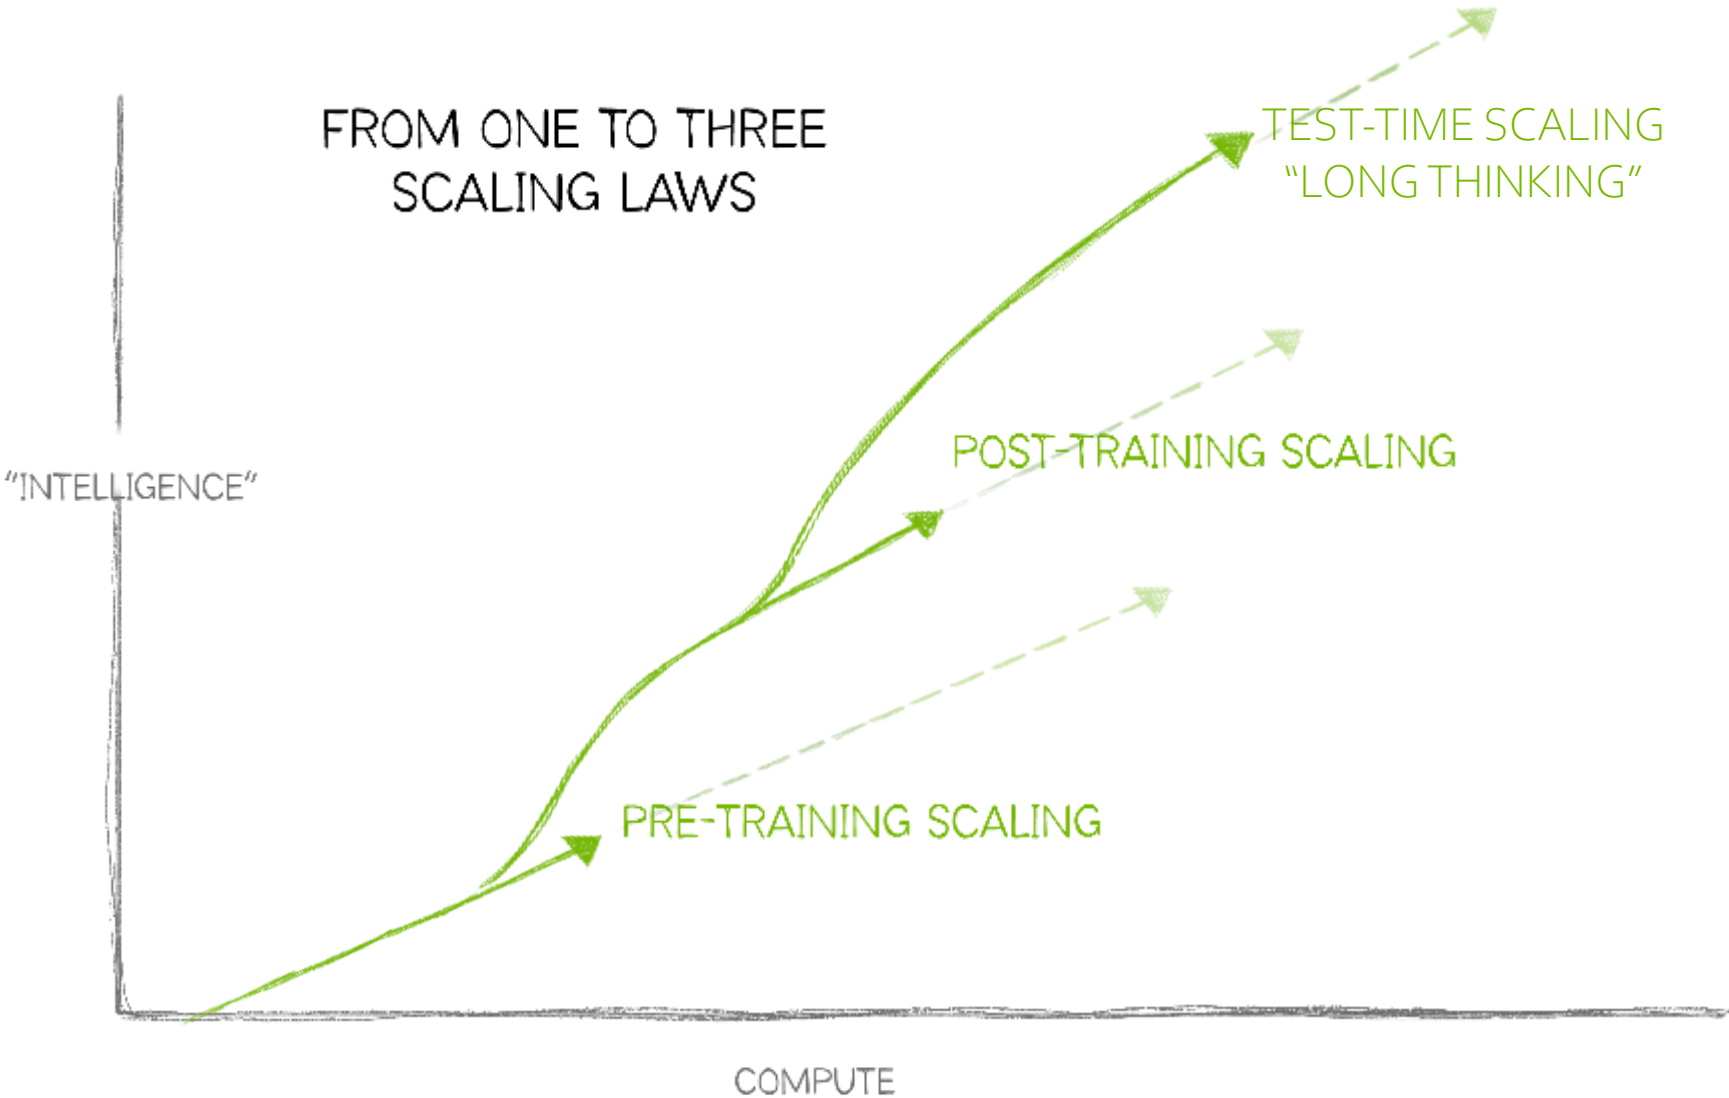


2012 ALEXNET

- Performance & Size of models doubling every 6 months
- 1 Million X improvement in 10 years!

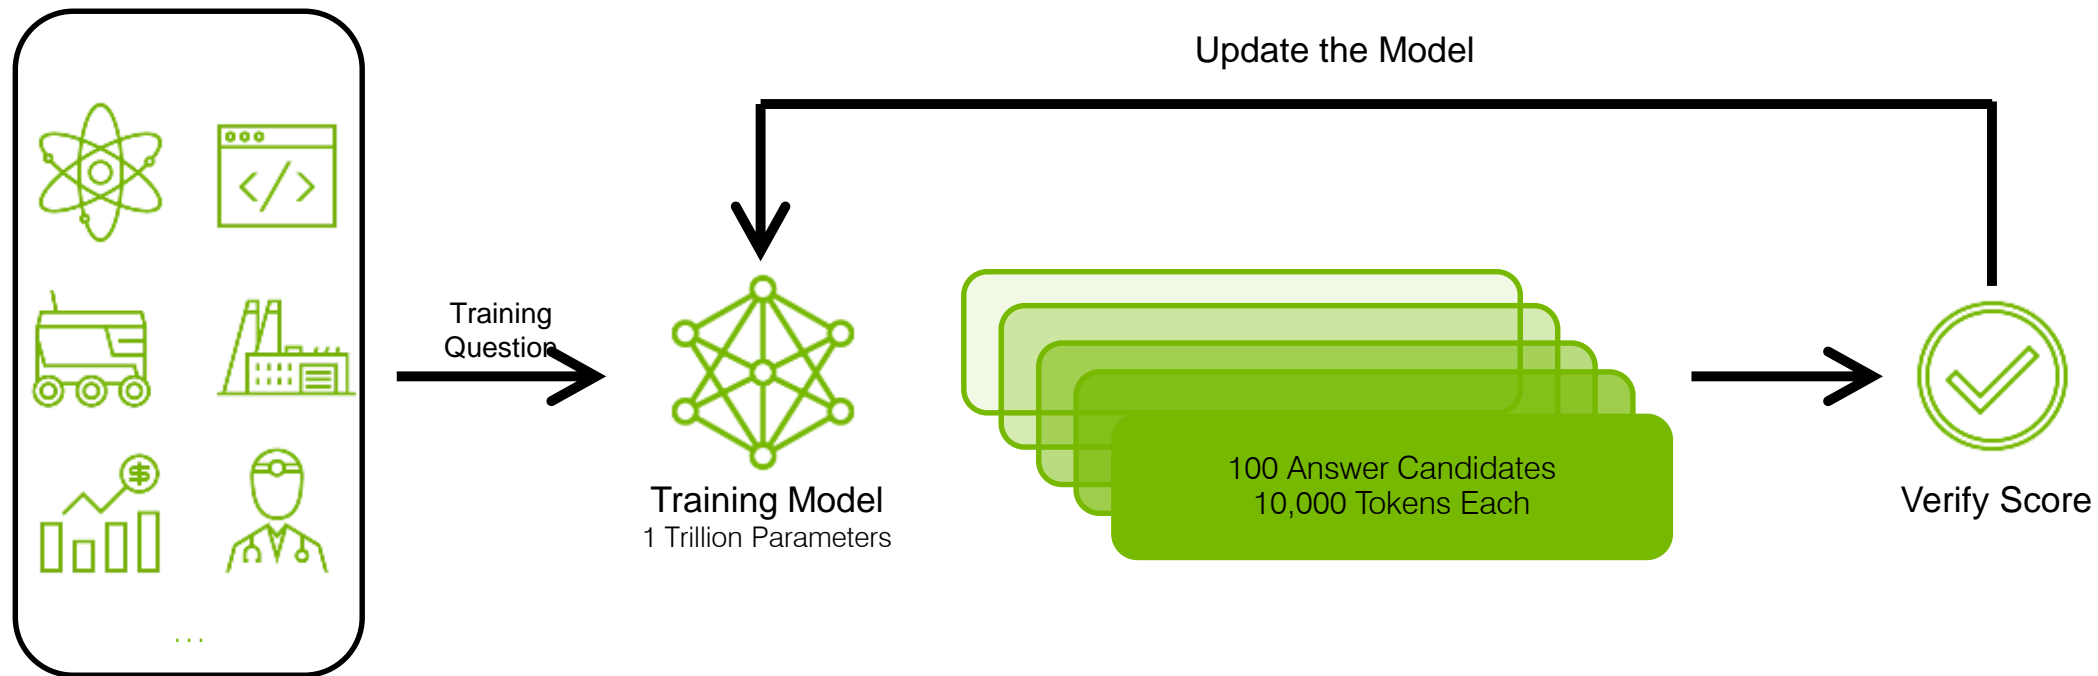
AI Scaling Laws Drive Exponential Demand for Compute

New "long thinking" supercharges inference scaling



Post-Training Scaling in an Era of Reasoning

Teaching models to think requires a lot of compute



100

Topic Areas

1 Million

Questions Per topic

100 Topics x 1 Million Questions x 100 Answers x 10,000 Tokens

100 Trillion Tokens

7X More Than What Llama 3 Was Pre-Trained On

Reasoning Drives 100X More Tokens and Revenue

Cost increasing with compute – requires full stack innovation to lower cost/token

KNOWLEDGE

Pre-trained models simulate internet knowledge
100 tokens | \$.50 / M Token

What is NVIDIA?

NVIDIA is a multinational technology company that specializes in designing and manufacturing graphics processing units (GPUs), high-performance computing hardware, and artificial intelligence (AI) technologies. The company was founded in 1993 by Jensen Huang, Chris Malachowsky, and Curtis Priem.

NVIDIA is headquartered in Santa Clara, California, and has become one of the leading companies in the field of computer graphics, gaming, and AI.

REASONING

Reasoning models think through problems
~10K tokens → 100X cost

Okay, so I need to figure out seating for a family dinner with 8 people. Let's list out the requirements first to make sure I don't miss anything. The user mentions: in-laws (which I assume are the wife's parents), parents (user's parents), sister, little brother, wife, and the user themselves. That adds up to 8: user, wife, user's mom, user's dad, wife's mom, wife's dad, sister, and little brother. Now, the constraints: in-laws (wife's parents) and user's parents (user's mom and dad) should not sit next to each other. The user's wife and user's mom should not sit next to each other. The user's sister needs to sit next to the user's little brother to help him eat. My wife really doesn't want to sit next to my mom. Where should each person be seated?

Additional context: The problem doesn't specify, so maybe I should assume a round table because it's common for family dinners and allows for a circular arrangement. If it's a round table, the seating problem might be easier with a circular arrangement.

To seat your family of 8 harmoniously at a round table while respecting all constraints, here's the optimal arrangement:

You Wife Dad Mom Sister Little Brother
Mother-in-law Father-in-law

Key Constraints Met:

- In-laws vs. Parents: Separated by your sister, brother, and you.
- Sister & Brother: Adjacent for assistance.
- Wife & Mom: Not next to each other (separated by you and dad).

Let's try:

1: S 2: LB 3: ? 4: ? 5: ? 6: ? 7: ? 8: ?

We need to fill positions 3-8 with U, W, UM, UD, WM, WD.

Let me consider placing the user (U) next to LB (position 2) to have a family member next to the little brother. So position 3: U. Then position 8 (next to S) could be someone else.

But then U at 3 would be next to LB at 2 and position 4. Let's see.

Alternatively, maybe place U and W opposite S and LB. Let's see, in an 8-seat round table, opposite of S (position 1) would be position 5. So position 5 could be U or W. Let's try placing U at 5 and W at 6, but need to check constraints.

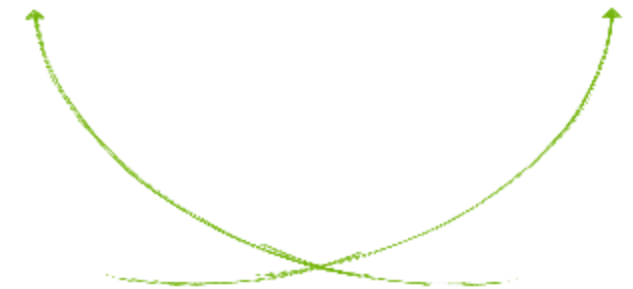
Wait, maybe a better approach is to divide the table into sections. Let's say S and LB are at 1 and 2. Then, to separate the in-laws from parents, we can place the parents (UM and UD) on one side, and in-laws (WM and WD) on the other side, with U and W in between.

ACCELERATING AI FACTORY VALUE

>\$1/M Token | >1T Param | >300 TPS

Need for increasing tokens

Increasing revenue customer experience



Decrease cost

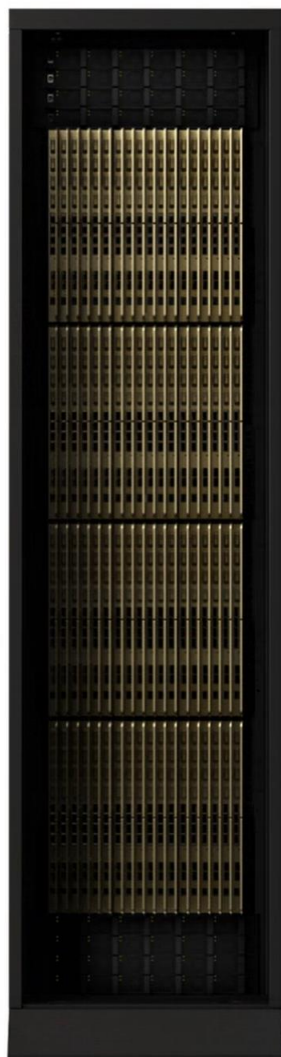
NVIDIA Paves Road to Gigawatt AI Factories

One-Year Rhythm | Full-Stack | One Architecture | CUDA Everywhere



Rubin Ultra NVL576

Flagship system for AI training and Inference workloads



Kyber Rack
Liquid Cooled

15 EF FP4 Inference
5 EF FP8 Training
14X GB300 NVL72

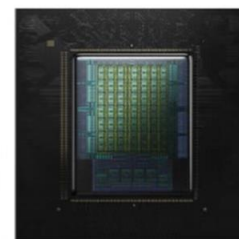
4.6 PB/s HBM4e
365 TB Fast Memory
8X

1.5 PB/s NVLink7
12X

1 15.2 TB/s CX9
8X

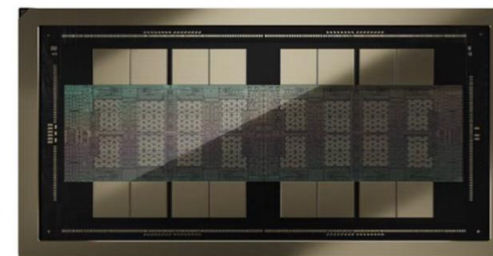
Availability
Second Half 2027

Vera



88 Custom Arm Cores
176 Threads
1.8 TB/s NVLink-C2C

Rubin Ultra



4 Reticle-Sized GPUs
100PF FP4 | 1TB HBM4e

THE RISE OF THE AI FACTORIES

European Union Building 20 AI Factories

Barcelona Supercomputing Center (Spain)

CINECA (Italy)

CSC (Finland)

GENCI/CEA (France)

GRNET (Greece)

Jülich Supercomputing Center (Germany)

Linköping University (Sweden)

LuxProvide (Luxembourg)

Poznan Supercomputing Center (Poland)

Sofia Tech Park (Bulgaria)

University of Maribor (Slovenia)

University of Stuttgart (Germany)

Vienna Scientific Cluster (Austria)

2 AI Factories (Location TBD)

5 Gigafactories (Location TBD)





EuroHPC
Joint Undertaking

New EuroHPC AI Factories



CZAI

1HealthAI



Gaia AI

LitAI



RO AI

NLAIF





EuroHPC
Joint Undertaking

EuroHPC AI Factory Antennas



AIFA ICE



Pharos-CY



HunAIFA



AIFA-LAT



AIF IRL



FAIMA



HEARTS

SKAIAT



CALYPSO



BE-AIFA



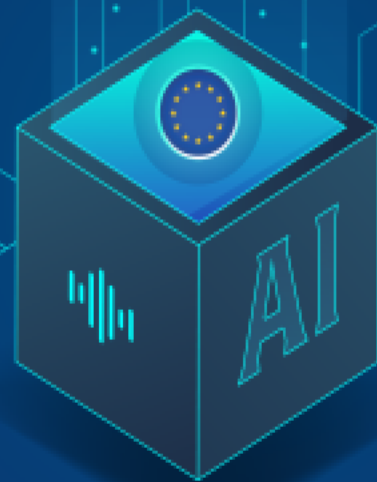
VEZILKA



UKAIFA



SAIFA



Stargate
200 MW, Scaling to 1.2 GW
Abilene, Texas



Courtesy of Crusoe



That's all Folks!

