

LARGE LANGUAGE MODEL HIGHLIGHT

Llama 2

Meta AI · USA · Llama 2 · Announced July 2023

TOTAL PARAMETERS

70B*Dense*

TRAINING TOKENS

2T*Web-scale*

DATA RATIO

29:1*Near-Chinchilla*


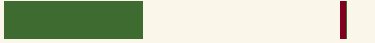
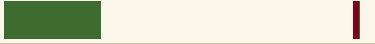
TRAINING COST

~\$1M*Hopper hardware basis***SUMMARY**

Llama 2 by Meta AI (USA) is a Dense model with 70B total parameters, trained on 2T tokens seen, for a training-token-to-total-param data ratio of 29:1, placing it in Near-Chinchilla territory.

On a normalized hardware basis (circa 2026 Hopper pricing), the estimated cost of a comparable new training run is ~US\$1M.

BENCHMARKS (VS LEADING MODEL IN RED)

MMLU		68.9
MMLU-Pro		37.5
GPQA		26.3
HLE	Not reported.	

ASSETS

Documentation <https://ai.meta.com/research/publications/llama-2-open-foundation-and-fine-tuned-chat-models/>

Playground <https://www.llama2.ai/>

More <https://lifearchitected.ai/models-table-pro/>

ADDITIONAL REMARKS

Context window=4096. MMLU=68.9 (GPT-3.5=70.0, GPT-4=86.4)

RATINGS (2026Q2)**61/100****COMPUTE PERCENTILE**

Percentile rank of estimated compute versus all models in Models Table Pro.

A**DISCLOSURE**

Parameters disclosed · Tokens disclosed · Doc available · Playground available.

-**ESTIMATE CONFIDENCE (*)**

Model data disclosed by lab, estimates not required.

C**OPTIMIZATION**

Tokens-to-total-params ratio: 29:1 · Near-Chinchilla.