

LLM

- <https://collabnix.com/best-ollama-models-in-2025-complete-performance-comparison/>

For Production Deployment:

- Primary Choice: DeepSeek-R1 32B for reasoning-heavy applications
- Coding Tasks: Qwen2.5-Coder 7B for optimal balance of capability and efficiency
- General Purpose: Llama 3.3 70B for maximum versatility
- Edge Computing: Phi-4 14B for resource-constrained environments

Optimization Strategies:

- Always enable **Flash Attention** and KV-cache quantization
- Use **Q4_K_M** quantization for production deployments
- Implement caching for repeated queries
- Monitor GPU memory usage and implement automatic model swapping
- Use load balancing for high-throughput applications


Hardware	Llama 3.3 8B (tokens/sec)	Llama 3.3 70B (tokens/sec)	Llama 3.2
RTX 4090	89.2	12.1	
RTX 3090	67.4	8.3	
A100 40GB	156.7	45.2	
M3 Max 128GB	34.8	4.2	
Strix Halo 128GB ollama		5.1	85.02
Strix Halo 128GB llama.cpp			90
RTX 3060			131.76

ROCM

model	capabilities	size	context	quantization	eval rate [token/s]	prompt eval rate [token/s]
qwen3-coder	completion tools	"30.5B"	262144	"Q4KM"	52.10	776.55
llama3.2	completion tools	"3.2B"	131072	"Q4_K_M"	52.78	1957.30
gpt-oss:20b	completion tools thinking	"20.9B"	131072	"MXFP4"	45.37	519.90

VULKAN

model	capabilities	size	context	quantization	eval rate [token/s]	prompt eval rate [token/s]
qwen3-coder	completion tools	"30.5B"	262144	"Q4KM"	54.03	805.43
llama3.2	completion tools	"3.2B"	131072	"Q4_K_M"	52.54	1838.82
gpt-oss:20b	completion tools thinking	"20.9B"	131072	"MXFP4"	43.36	475.60



From:

<https://wiki.csgalileo.org/> - **Galileo Labs**

Permanent link:

<https://wiki.csgalileo.org/tips/llm?rev=1771053520>

Last update: **2026/02/14 08:18**

