

Soup: Fine-Tuning an 8B Model on a 4 GB Laptop GPU, One Layer at a Time

Kabui, Charles

2026-09-14

[Read at ToKnow.ai](#)



Soup is an open-source tool that fine-tunes large language models from a single YAML config. Its headline feature, [layer streaming](#), keeps the frozen base out of GPU memory: the weights stay in CPU RAM and are copied to the card one layer at a time, double-buffered so the next layer loads while the current computes. Quantizing that base to 4-bit shrinks it about fourfold. Llama-3.1-8B then fine-tunes at 119.6 tokens per second with a 3.32 GB peak on a 4 GB RTX

3050 Laptop, a card that cannot hold the model. A Qwen2.5-3B run with an unquantized 16-bit base reaches 143 tokens per second in 2.15 GB, a setup that runs out of memory trained resident.

LoRA fine-tuning has meant renting a cloud GPU or a bigger card; here a 4 GB laptop works, about 1.43 times slower than a resident run. The project is upfront that layer streaming is not new and that the published small-card result is 1.3B, from [LSP-Offload](#) without streaming. Soup's narrower claim is up to 8B by streaming the base, with the forward pass verified bit-exact against a resident run. If the model already fits your card, leave it off.

In August the project documented a bug where streamed 4-bit runs above roughly 165 MB per layer produced silently wrong gradients at 32B and 72B while the loss curve looked healthy; v0.73.0 fixed it. The 119.6 tokens per second figure predates that repair and has not been re-measured on that card since.

Read More: [Running a frontier model in 18 GB of RAM](#).

Sources:

- [Soup CLI on GitHub](#)
- [Soup docs: Layer streaming, train past your VRAM](#)
- [Exact Layer Streaming: LoRA Fine-Tuning of an 8B Model on a 4 GB Laptop GPU \(Zenodo\)](#)
- [Soup docs: Validation on hardware we do not own](#)
- [Practical offloading for fine-tuning LLM on commodity GPU via learned sparse projectors \(arXiv\)](#)

Disclaimer: For information only. Accuracy or completeness not guaranteed. Illegal use prohibited. Not professional advice or solicitation. **Read more:** [/terms-of-service](#)