Episode Book

EPISODE #6228

What 12GB of VRAM Can Actually Do Now

PUBLISHED	23 October 2026 at 13:28 UTC
DURATION	22:31
LISTEN	myweirdprompts.com/episode/12gb-vram-local-ai-guide-87609/
HOSTS	Corn, Herman, Daniel, Hilbert
SCRIPT MODEL	anthropic/claude-sonnet-5.5
VOICES	chatterbox-regular (chatterbox-local)
PIPELINE	V5.20
BOOK GENERATED	24 October 2026 at 07:26 UTC

SYNOPSIS

A Radeon RX 7700 XT owner gave up on running models locally six months ago — the context window filled instantly and the model lost the plot. But the models changed underneath him. This episode walks through the new memory arithmetic: why a 35B mixture-of-experts model with only 3B active can run on a 12GB card, how cache costs dropped roughly seven-fold between model generations, and what that means for speed, fine-tuning with Unsloth's official AMD support, and the honest limits nobody has benchmarked yet.

TRANSCRIPT



Corn

Daniel's last serious try at running a model on his own desktop was about six months ago, and his verdict was short. The context can't hold anything, so what's the point. He's got an AMD Radeon seven seven hundred XT in there, twelve gigabytes, which is a perfectly respectable gaming card and apparently not a respectable AI card.



Herman

Which is the interesting part, because the verdict was fair when he reached it and it's mostly out of date now. The models changed underneath him. Something shifted in how they remember, and it shifted a lot.



Corn

He wants to know what a normal person with a mid-range card can actually do today. Fine-tune, transcribe, clone a voice, call tools. And whether he should be pleasantly surprised or just quietly relieved he didn't spend more.



Daniel

I'll be honest, I've mostly given up. I tried a few things, the context filled up almost immediately, the model lost the plot, and I thought, I can just use an API for this. If you've got something that changes my mind, go ahead.



Herman

Let's do the money first, because it frames everything. That card is around four hundred dollars new in the United States, give or take, and that number comes from a fairly stale price tracker, so treat it as a ballpark. An NVIDIA card with sixteen gigabytes, an RTX fifty seventy Ti, has been selling somewhere between nine hundred fifty and twelve hundred fifty depending on whose listing you believe.



Corn

So he paid a third of the price for three quarters of the memory.



Herman

And if you want twenty-four gigabytes from NVIDIA, a used three thousand ninety Ti was listed around fifteen hundred in July. Roughly four times the card on his desk. That's why twelve gigabytes on a Radeon is where a lot of normal people actually are. The price gap is also the reason AMD owners spent two years being told, politely, to buy a different card.



Daniel

Okay, but twelve gigabytes was exactly the problem. Eight billion parameters, a few thousand tokens of context, and it's full.



Herman

That memory of yours is accurate for the twenty twenty-five generation. Take an older design like Qwen three eight B. By my arithmetic, and I'm working from a config I recall rather than one I've checked, every token of context costs about a hundred and forty-four kilobytes of cache. Thirty-two K of context is four and a half gigabytes. A hundred and twenty-eight K is eighteen gigabytes, which doesn't fit on your card at all.



Corn

Eight B at thirty-two K and the cache alone eats the card.



Herman

Now the new generation. Qwen three point six, the thirty-five B with three B active, has forty layers but only ten of them do full attention. The rest use a cheaper recurrent mechanism that carries a small fixed state. Those ten layers have two key-value heads each. So the cost is about twenty kilobytes per token, in the full sixteen-bit format.



Corn

Twenty versus a hundred and forty-four. Seven times smaller.



Herman

Thirty-two K of context is about six tenths of a gigabyte. A hundred and twenty-eight K is two and a half. Run the cache in eight-bit and halve that, with very little quality loss in most reports. Even the full native context of two hundred sixty-two thousand tokens is around five gigabytes in sixteen-bit.



Daniel

Hang on. That's the cache. The model itself is still thirty-five billion parameters. That's not going on a twelve gigabyte card.



Herman

Not all of it, no. The Q four weights are around twenty-two gigabytes. This is the trick I want to get to, and you've just set it up better than I would have.



Corn

Before you do, I want to say that fitting the context is not the same as the model using it well. Two hundred sixty thousand tokens in memory and a model that gets lost at forty thousand is still a model that gets lost.



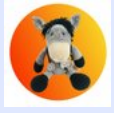
Herman

Fair, and I can't tell you how well it holds up at a hundred K on a card like yours. Nobody's put a number on that for this setup. What I can say is that the reason you hit a wall was memory, and the memory arithmetic is now completely different.



Corn

You're saying his problem got solved by a spreadsheet.



Herman

By a change in the layer layout, yes. Gemma four twelve B does a version of the same thing, with mostly sliding-window layers of a thousand tokens and only occasional global ones. I don't have a clean number for its cache, so I won't invent one. One practical



Herman (continued)

warning on the eight-bit cache. It needs flash attention switched on, and on some NVIDIA builds a flag can silently push attention onto the CPU and make prompt processing crawl, by a factor of twenty-odd. I don't know if the Vulkan or ROCm builds have the same trap. So if reading a long prompt suddenly takes forever, that's the first thing I'd check.



Corn

Fine, but it's still thirty-five billion parameters. Nobody's squeezing that onto twelve gigabytes.



Herman

A mixture-of-experts model has lots of experts but only runs a few per token. Three billion parameters are active at a time. So you keep the attention layers and the shared parts on the GPU, and you park the expert weights in system RAM. There's a flag in llama.cpp, the n-cpu-moe option, that says how many layers' worth of experts stay on the CPU.



Corn

That sounds like it should be slow. Everything crawling through the memory bus.



Herman

You'd think so. One report on an RTX three thousand sixty, twelve gigabytes, same size of card as yours, ran the Qwen three point six thirty-five B at a four-bit-ish quantisation with eighteen layers of experts offloaded. About nine hundred tokens a second reading the prompt and about forty-seven generating. That's an anecdote on NVIDIA, a card with less bandwidth than yours.



Corn

Forty-seven tokens a second is faster than I read. Which, to be fair, is a low bar.



Herman

Quite a bit faster than the average person talks. On the Radeon, I'd expect lower, possibly meaningfully lower, and that's a guess by the maths, not a measurement. But the shape holds. And what decides it is system RAM. Twenty-two gigabytes of weights has to live somewhere. Daniel, do you know what's in the box?



Daniel

I'd have to check. I'd rather look than guess. Is that the thing that decides it?



Herman

It's the main thing. Sixteen gigabytes free would be tight, thirty-two is comfortable. And it changes the old advice. A nine B dense model uses all nine billion parameters on every token. The thirty-five B mixture uses three. On a card like yours the bigger model can be the faster one, as long as the RAM holds. Corn, which would you pick, the small one that fits entirely, or the big one leaning on the CPU?



Corn

The small one for anything I need before lunch. The big one for anything I'm willing to leave running while I nap. Which is most things.



Herman

Fair. Speed in general, then, since I know you'll want it. For a dense eight B model at four-bit, one fit-calculator puts that card around fifty-seven tokens a second. That's an estimate, not a benchmark. A neighbouring card, the seven eight hundred XT, was measured at roughly a hundred and twenty tokens a second on an older seven B model in Vulkan.



Corn

And the ceiling is memory bandwidth.



Herman

Four hundred thirty-two gigabytes a second on yours. Every token reads the whole active weights once, so bandwidth divided by model size is your upper limit. Which is also why the three-B-active model is so quick. It only reads three billion parameters' worth per token.



Daniel

What about ROCm? I'm still on six point four. Do I have to upgrade?



Herman

For chat, probably not. In the tests I know of, Vulkan generates tokens faster than ROCm in llama.cpp, and ROCm wins on prompt processing and time to first token. Vulkan needs no ROCm at all. The upgrade matters more for training and for anything PyTorch-based. We did a whole episode on ROCm ten point one growing up, and the short version is that the software side really did get better.



Corn

So he can stay on six point four, run Vulkan, and nobody calls the ROCm police.



Herman

For chat, yes. The dense models people actually recommend for this tier are Qwen three point five nine B, which is around six and a half gigabytes at four-bit, and Gemma four twelve B, which one source puts at six and a half and another at eight. I'm not sure which is right, and the difference decides how much context you can afford next to it. Granite and Phi get named too.



Corn

Six and a half gigabytes leaves five and a half free on his card. That's a lot of room for context.



Herman

With the cache sizes we just did, a very comfortable amount. Which is the whole change in one sentence: the weights used to be the cost, and now the weights and the cache both fit. For the other stuff I'd wait and see what each tool asks for. Which brings up the one I know you've been waiting for. Fine-tuning.



Corn

Unsloth. I'm a little embarrassed to say I know how it works, because I founded it.



Herman

You didn't found Unsloth.



Corn

No. I was actually a technical consultant.



Herman

To Unsloth.



Corn

For eleven years. Which was a long time, mostly because of how I worked. I would arrive on a Monday and be fully present by about Thursday. I advised on the logo. That's the sloth in the logo, so that was largely me.



Herman

Unsloth was not founded by a sloth.



Corn

They weren't there for the early years, Herman. I had a desk. It was a very low desk, and I advised from it, slowly, for what I'd call the first formative stretch. My main contribution was a memo on the word "faster". I recommended it. It took me six months to write the memo, and I was told it was the shortest memo ever submitted.



Herman

It was one word.



Corn

Two. I had a draft with a hyphen. I was also in charge of a plant that stood in the lobby. I watered it on a rotation of my own invention, once per season, and it thrived, and I think the team felt that.



Herman

Corn.



Corn

Then there was the offsite, which I attended in spirit, and in a hammock, in what I believe was Mongolia...



Herman

Okay, stop. The actual question is whether twelve gigabytes can fine-tune anything.



Corn

I was getting to the best part, which is that I never invoiced them.



Herman

That tracks. Unsloth added official AMD support in July. The RX seven thousand series, including this card's chip, gfx eleven oh one, is listed as full support on Linux, Windows and WSL. The installer picks the right ROCm PyTorch, bitsandbytes and llama.cpp builds for you, so you're not hand-assembling the stack.



Daniel

What size can I actually train? I assumed a twelve gigabyte card meant tiny models.



Herman

Unsloth publishes minimums for QLoRA at four-bit. Three B needs about three and a half gigabytes. Seven B is five. Eight B is six. Nine B is six and a half. Fourteen B is eight and a half, and a twenty-seven B is twenty-two, which is out.



Corn

So he can do anything up to nine comfortably and fourteen if he's brave.



Herman

Those are absolute minimums, so once you add activations and context, fourteen gets marginal and you'd drop the batch size to one. Nine B, I'd call comfortable. But I'll be straight about this: I haven't found a single published run on a seven seven hundred XT or a seven eight hundred XT. The support is official. The evidence for this exact card doesn't exist yet, and one third-party guide warns that the stack is fragile when kernel, driver and PyTorch don't line up.



Corn

What does he even get out of it, though? A nine B model that's learned his writing style?



Herman

Style, format, house vocabulary, a narrow task done the same way every time. QLoRA is not how you teach a model new facts, and people get that wrong constantly. It's how you teach it manners. For that, nine billion parameters is plenty.



Corn

Daniel could be the first data point.



Herman

He could. Corn, honest question, would you fine-tune on a card you can't find a single report for?



Corn

I'd start with a three B model and see if the machine survives. And I'd leave the room.



Herman

On twelve gigabytes you can fine-tune a model, if you're patient. Slowly, but you can. Ivan: Larkspur. Seven, seven. Two, nine. Four, one. Seven, seven. The rye is in the proving basket. The baker counts his own fingers. The twelve gigabyte loaf is not for the oven. Apricot.



Herman

Gentian. Three, three. The ferry runs late on Thursdays. Quince.



Corn

What on earth was that?



Herman

Speech recognition next. Fine-tuning Whisper is lighter than fine-tuning a language model. Whisper large v three turbo is the large model with the decoder cut from thirty-two layers to four. One vendor calculator puts a LoRA run at about two and a half gigabytes in sixteen-bit, and a full fine-tune around seven. That's modelled at batch size one, so it's a floor.



Daniel

Fine, but what about just using it? I do a lot of dictation.



Corn

What does he gain from fine-tuning a transcription model? Whisper already handles most speakers.



Herman

Names and jargon, mostly. Whisper writes a company name, a library, a person's surname however it sounds, and it gets the same ones wrong every time. A small LoRA on a few hours of your own recordings with a corrected transcript fixes exactly those. It's a lot of value for a model that fits in a few gigabytes.



Corn

A few gigabytes. On a twelve gigabyte card that's practically a rounding error.



Herman

Inference is where the trouble is. The popular fast engine, faster-whisper, sits on CTranslate two, and the version you install from the package index is CUDA-only. On a Radeon it quietly falls back to the CPU. There are community forks that run it on ROCm, but they're niche.



Corn

Of course the fast one is the one that doesn't work.



Herman

The way out is whisper.cpp with the Vulkan backend. Cross-vendor, packaged on Arch, no ROCm required. I haven't seen a number on a seven thousand series card. On an AMD integrated chip, the large model ran about four and a half times faster than real time, which is the wrong hardware but gives you a feel. We covered local speech to text for AMD owners in an early episode of the show, and the news is that the escape hatch is better paved now.



Corn

Dictation doesn't need a jet engine.



Herman

Then a discrete card with this much bandwidth should beat it comfortably. By the maths, anyway. Now Chatterbox, which, full disclosure, is the engine reading this episode.



Corn

We're asking whether a twelve gigabyte Radeon can run what's running us.



Herman

The model's about half a billion parameters, so memory is not the issue. What I found is a single account from a Strix Halo machine, an AMD chip with integrated graphics. About ten seconds to generate a clip, working from a ten-second voice sample. It needed a nightly PyTorch build, because the stock ROCm six point four wheels didn't contain kernels for that chip and it fell back to the CPU.



Daniel

So that's me. Six point four.



Herman

Possibly not. Your chip, gfx eleven oh one, is a mainstream target, so it shouldn't hit that exact wall. That's an inference, not something I've seen confirmed. In production the show runs at roughly five times faster than real time on an NVIDIA L four. I'd treat matching that as the goal and expect to fiddle on the way. This is also the one place where I'd consider upgrading ROCm first.



Corn

So the summary is: probably works, nobody's written down that it does.



Herman

The model's small enough that the fit is a non-issue, so what you're really testing is whether the software stack builds cleanly on your machine. If it does, a voice clone from a short sample is well within reach. If it doesn't, you lose an evening to wheels and driver versions.



Corn

Which is the real price of the cheaper card. You pay with evenings.



Herman

Hm. Fair.



Corn

Now tools. A model that invents arguments, or calls the right tool forever, is what I'd expect to hear about.



Herman

Both of those still happen, and I'd rather concede it now than have you find out. For a single tool with clear parameters, the current models do well. Qwen three point six's own model card says it excels at tool calling, and it's served with a dedicated parser. Multi-step flows, where step three depends on a subtle reading of step two, are where local models still lag Claude or GPT or Gemini, and I haven't seen a clean head-to-head benchmark to put a gap on it.



Corn

So it's a good intern. Not a good colleague.



Herman

Corn, would you hand an intern the keys to your calendar and your email?



Corn

Read-only. And I'd check it Thursday.



Herman

That's a perfectly good rule for a twelve gigabyte model, actually. Give it a read-only folder, one tool at a time. Then you see where it breaks. The plumbing matters too. LM Studio is an MCP host itself, so you add a server to a config file and it works. Ollama does tool calling but isn't an MCP client, so you need another layer on top. And quantisation matters: the four-bit K M variant is the practical floor. Go lower and the tool calls get sloppy first.



Corn

I'd still reach for the API for anything that has to be right.



Herman

That's reasonable. Use the local model for the things where thirty-two K of context and a slow Tuesday afternoon are plenty.



Hilbert

K isn't a thousand, Herman. Thirty-two K is thirty-two thousand, seven hundred and sixty-eight. You keep rounding it down. I've got a Mr. Leszek Brandvold at the Regional Metrology Office stamping my certificate at thirty-two thousand flat. He says K is a thousand. I'm seven hundred and sixty-eight tokens short and he's holding it until I send the difference. Nineteen days. I've mailed two hundred and eleven so far. One per postcard.



Herman

He's right, Hilbert. In context lengths, K is a binary thousand. We've been sloppy and it makes no difference to anything we've said, but he's right.



Corn

Two hundred eleven in nineteen days is about eleven a day. He's got another fifty to go.



Hilbert

Postage went up on Monday.



Herman

The missing tokens cost very little. At twenty kilobytes each on that Qwen layout, seven hundred sixty-eight of them is about sixteen megabytes of cache. Mr. Brandvold could be satisfied in memory terms by a rounding error.



Corn

He'd still want the postcards.



Hilbert

He takes paper only. I'll mention the sixteen megabytes.



Herman

Daniel, so what would I actually do this weekend, on that card? Saturday morning, nothing exotic. Put llama.cpp or LM Studio on it with the Vulkan backend, leave ROCm at six point four, and pull one of the dense nine to twelve B models: Qwen three point five nine B or Gemma four twelve B, around six and a half to eight gigabytes at four-bit. See what the speed feels like. My estimate is around fifty tokens a second, and it's an estimate.



Corn

Then check the RAM, and if it's thirty-two, try the big one with the experts parked in memory. Raise the context to thirty-two thousand seven hundred sixty-eight and watch what it does to the memory graph. I'd like to see that the cache really is as small as the maths says.



Herman

And if the RAM is sixteen, it'll be tight on that model. Stay with the dense ones, and treat the thirty-five B as a thing to try when you add RAM. RAM is the cheap upgrade in all of this.



Corn

Sunday I'd do whisper.cpp with Vulkan on a recording of yours and see if it keeps up with your talking.



Herman

Which is a dare, with Daniel's speaking speed. After that, if you're feeling brave, Unsloth's installer and a small QLoRA run on a three B model. Not a nine B first. And Daniel, if it works, nobody's published a seven seven hundred XT run. You'd be writing the post.



Daniel

I might do that. Chatterbox I'd save for the weekend after. I'll fix the ROCm version when it tells me to.



Corn

Common belief, then: local means small and dumb. A twelve gigabyte card means a toy.



Herman

It meant that when memory was the bottleneck for everything. A thirty-five B model that only activates three B per token, with the experts in RAM and a cache of under a gigabyte at thirty-two K, is not a toy. It's slower than an API and weaker at multi-step tool work. But it's not a toy.



Corn

What I'd like to see is the real figures on the cards people actually own. Sixteen megabytes of cache is a rounding error on paper. Whether the whole stack behaves on a seven seven hundred XT is still a postcard nobody has sent.



Herman

Thanks to Hilbert Flumingtop, our producer, for keeping the show running.



Corn

This has been My Weird Prompts. If you liked it, the website is my weird prompts dot com, and there's more there.



Herman

See you tomorrow.



Corn

See you tomorrow.