AI YouTube Channels to Follow in 2026: 185 Ranked by Signal

A data-backed directory of roughly 185 AI YouTube channels across 24 active thematic tiers, ranked by recent reach, recency, signal quality, and verification rather than subscriber count.

Yiğit Konur · Jul 13, 2026 · 56 min read

▶ The useful bits

Snapshot	Notes
Coverage	Roughly `185 channels` across `24 active thematic tiers`, plus a legacy tier, graveyard, blacklist, and not-admitted log
Primary signals	User-supplied live scrape from `2026-07-12/13`, recent views over approximately the latest 15 public uploads, three research synthesis passes, channel activity, and attributed Reddit reputation signals
Bottom line	Rank a feed by `recent reach × recency × signal quality × verification`, not subscriber count; mix evergreen curricula with a small active frontier and work-specific stack

Executive Summary

Choosing who to follow in AI YouTube in 2026 feels like navigating a minefield. There are eight-million-subscriber legends who have not uploaded in a year, a 763K-subscriber channel pulling roughly 2.7K views per video, and a wall of daily "SHOCKING!!! AGI IS HERE!!!" thumbnails engineered to eat an evening. Meanwhile, a useful feed has to make decisions quickly.

This report starts from a user-supplied research corpus that live-scraped 68 core channels across **2026-07-12/13**, merged those results with three deep-research passes, reconciled conflicting stats channel by channel, de-duplicated alias collisions su[Skip to content]ser" versus "Token Chasers," and expanded the set into roughly **185 distinct channels across 24 active thematic tiers**. A separately labeled legacy tier sits inside the graveyard section and is preserved rather than renumbered away.

This is not a “top 10 AI YouTubers” list. It separates recency from signal quality and verification, preserves Reddit warnings alongside positive recommendations, distinguishes live feeds from evergreen reference libraries, and ends with curated stacks for foundations, research awareness, daily news, local models, agentic engineering, and image/video systems.

The most important result is the gap between subscriber badges and current reach. Fahd Mirza has 763K subscribers and averages roughly 2.7K recent views; Anthropic has 732K subscribers and averages roughly 315K. Similar badges, about **115×** difference in recent reach. Depth can also scale: Dwarkesh Patel averages roughly 309K on long technical interviews, while bycloud averages roughly 83K on architecture explainers.

Every minimum, maximum, and average view figure in this report is computed over approximately each channel’s **latest 15 public uploads**, not lifetime performance. Dormant channels can therefore show large averages carried by evergreen hits, while daily channels reveal their current pull. Karpathy’s roughly 2.06M average is the gravity of a remarkable back catalog, not evidence of a current publishing feed.

> ***Provenance note:*** *The underlying dataset was supplied by the user and produced through guided deep research plus live scraping. I transformed it into this report structure, preserved its honesty flags, and did not independently re-scrape or silently replace missing data. Every ❓ remains a claim the source corpus could not verify.*

How to Use This Report

The report is intentionally long. Use the route that matches the decision in front of you.

Your situation	Start here	Typical time
“I want to understand the math.”	Part 1: The Foundations	20 min
“Keep me research-aware without living on arXiv.”	Part 2: The Frontier	15 min
“Te	Part 3: The News Diet	10 min
“I am building a local LLM rig.	Part 4: The Hardware Truth	15 min
“I ship code with agents daily.”	Part 5: The Builders	15 min

Skip to content

Your situation	Start here	Typical time
“I make images, video, or weird programs.”	Part 6: The Creators	10 min
“Who should I unfollow?”	Part 7: The Graveyard and the Blacklist	5 min
“Just give me the stacks.”	Part 8: Subscription Strategy	5 min

Evaluation Framework

Every channel carries three grades. They answer different questions, and collapsing them into one score is how a feed fills with channels nobody watches.

Recency — is this a live feed or a museum?

- 🔥 posted this week
- ✅ posted this month
- 🐢 posted within one to three months
- 💤 four months or more, effectively dormant as a feed

Signal quality — what does the content deliver?

- 🟢 high-signal, deep, or low-hype
- 🟡 solid and practical, but mixed
- 🔴 hype-flagged or contested in the supplied Reddit synthesis

Verification — did the data survive scraping?

- ❓ the channel page failed across Jina, Scrape.do, and Kernel during this pass; the report flags those metrics instead of guessing

Why Subscriber Counts Lie

The subscriber count is an archaeological artifact: it says a channel was worth subscribing to at some point. The recent-view average is a stronger signal of whether it de feed now. The comparisons below therefore emphasize **recent-vie ency × signal grade**, with verification status carried separately.

Skip to content

Stat Conflicts and Reconciliation

The supplied research passes occasionally disagreed. The freshest scrape is primary; discrepancies are disclosed rather than silently averaged.

Channel	Pass A figure	Pass B figure	Resolution
Welch Labs	400K subs	887K subs	**887K**; fresher scrape, and `@WelchLabs` / `@welchlabs` are the same channel
Julia Turc	48K subs	72.1K subs	**72.1K**; fresher, one channel at `@juliaturc1`
Token Chaser	25K	7.5K	**7.5K**; fresher live scrape, and "Token Chaser" / "Token Chasers" are the same creator
Donato Capitella	8K subs	99.7K subs	**99.7K**; the 8K figure was stale
Nate B Jones	6K subs	305K subs	**305K**; the 6K figure was a mis-scrape
GosuCoder	15K	27.1K	**27.1K**; fresher scrape
Emergent Garden	200K	273K	**273K**; fresher scrape
Asianometry	941K	944K	**944K**; fresher scrape, trivial growth delta
Brian Casel	23K	69.5K	**69.5K**; fresher scrape
Latent Vision	unverified ❓	38.5K confirmed	**38.5K verified**; upgraded to 🟢✅

Latent Vision is the most useful correction in the table: the core scrape could not verify it, while a later pass confirmed 38.5K subscribers and recent view counts. That changes its position in the ComfyUI tier materially.

Part 1: The Foundations

Math, statistics, from-scratch implementation, and structured courses. This is where you go to understand AI beneath the abstraction layer — and it's the part of YouTube where dormancy matters least, because the content is evergreen.

Tier 1 — Mathematical Foundations & Visual Intuition

The "understand it beneath the abstraction layer" tier. Raw math, physics, low-level model mechanics. Slow cadence, evergreen value.

Ch			Recent views (min–max)	Avg	Latest-upload themes	Gra de
3B wn	own	M	403K–6.1M	~1.69M	Entropy / "compression is intelligence", Laplace transforms, topology, quantum computing, "how AI images actually work"	🟢 ✅

Skip to content

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
StatQuest	@statquest	1.66M	1.8K–12K*	~6K*	Linear programming / Simplex, Random Forests pt.2, False Discovery Rates	🟢🔥
Welch Labs	@WelchLabs	887K	~297K (recent single); historic 12K–680K	~130K–297K	“What the Books Get Wrong About AI [Double Descent]”, *Neural Networks Demystified*, backprop/matrix visual series	🟢🐢
Steve Brunton	@Eigensteve	540K	5.8K–51K	~14K	Optimization bootcamp, Bayesian regression, Monte Carlo, hypothesis testing	🟢🔥
ritvikmath	@ritvikmath	211K	1.8K–14K	~5.1K	Regularization, conjugate priors, contextual bandits, the 4 must-know LLM params	🟢🐢
Serrano Academy	@SerranoAcademy	194K	912–13K	~4.2K	Neural nets bending space, RAG / vector DBs, tokenization, GRPO / DeepSeek	🟢🔥
Julia Turc	@juliaturc1	72.1K	25.5K–36.3K	~30K	World models, flow-matching physics, MoE gating, tensor cores, FP4 quantization, efficient serving	🟢✅

**StatQuest partial — only 3 recent view counts rendered during the scrape.*

3Blue1Brown — If you subscribe to exactly one channel from this entire directory, it’s this one. Grant Sanderson’s recent run — entropy and “compression is intelligence,” Laplace transforms, how AI images actually work — is the rare feed where a twenty-minute video genuinely replaces a semester. 8.47M subscribers and a ~1.69M recent average means this is that near-impossible thing: a channel whose reach *matches* its badge.

StatQuest — Still uploading weekly (Simplex method, Random Forests, False Discovery Rates) while everyone assumes he’s a legacy act. Here’s the thing about those tiny recent view counts: statistics fundamentals don’t trend, but they compound. This is the channel you’ll be grateful for when your eval pipeline produces a p-value you can’t explain.

Welch Labs — The legendary *Neural Networks Demystified* channel, repeatedly praised in r/learnmachinelearning for animating backpropagation and matrix operations step-by-step. Note the reconciliation: one research pass had him at 400K subs; the fresh scrape says **887K**. The recent double-descent video (“What the Books Get Wrong About AI”) is exactly the kind of content that justifies the badge — he tells yoSkip to contentooks are wrong, with animations to prove it.

Steve Brunton — A University of Washington professor running what is essentially a free graduate program: optimization bootcamps, Bayesian regression, Monte Carlo methods. Posting 🔥 this week. If your linear algebra is solid but your *applied* math is shaky, this is your gap-filler.

ritvikmath — Compact, whiteboard-style treatments of the topics interviews actually probe: regularization, conjugate priors, contextual bandits. The "4 must-know LLM params" video is a good litmus test for whether his density-per-minute works for you.

Serrano Academy — Luis Serrano has a gift for geometric explanations — "neural nets bending space" is exactly what it sounds like, and it sticks. His recent GRPO/DeepSeek coverage shows he's keeping the visual-intuition method pointed at current architectures, not just classics.

Julia Turc — The sleeper pick of this entire tier, and possibly this entire guide. Ex-Google Research engineer (YC S24), widely regarded as one of the best technical resources on YouTube for high-level mathematical explanations of *modern* architectures: MoE gating, flow-matching physics, tensor cores, FP4 quantization. 72.1K subs with a ~30K average — that ratio (over 40% of her subscriber base watching every video) is the highest engagement signal in this tier. Cross-listed in the local-inference tier for her serving/quantization work.

Tier 1.5 — Statistics & ML Prerequisites

Not AI-first feeds — but materially useful for the foundations behind ML evaluation, regression, probability, and experimental reasoning. Treat these as supplements, not subscriptions.

Channel	Handle	Subs	Recent views	Avg	Focus	Grade
Brandon Fol	@BrandonF	—	—	—	Intro statistics, finite math, management science, statistical learning	🟡💤 (uploads ~2 yrs old)
zec cs		—	—	—	Regression, survival analysis, distributions, hypothesis testing, probability intuition	🟡💤 (uploads ~3 yrs old)

Skip to content

Channel	Handle	Subs	Recent views	Avg	Focus	Grade
Reducible	@Reducible	—	—	—	Animated CS concepts — A* search, Fourier, TSP, PageRank, image compression	🟡💤 (newest ~1 yr ago)

All three are dormant, and it doesn't matter. **zedstatistics** got called "the best one" in a data-science thread for probability intuition, and probability intuition hasn't changed since 2023. **Reducible** is the closest thing to "3Blue1Brown for computer science" — A* search, Fourier transforms, PageRank — with a stated (never fulfilled) intent to extend into ML. Use them like textbooks: pull the chapter you need.

Tier 2 — LLM-from-Scratch & Deep Implementation Engineering

Code the transformer / GPT / attention / training loop by hand. This is the developer's canon — and here's the uncomfortable truth up front: **most of it is dormant.** The averages below are inflated by evergreen hits. Treat these as reference libraries, not fresh feeds.

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
Andrej Karpathy	@AndrejKarpathy	1.57M	36K–8.1M	~2.06M	"How I use LLMs", reproduce GPT-2, build the GPT tokenizer, makemore, micrograd	🟢💤
Sebastian Raschka	@SebastianRaschka	90K	11K–168K	~52K	LLM architecture comparisons, *Build-an-LLM-from-Scratch* series, finetuning, pretraining	🟢🐢
Umar Jamil	@umarjamilai	85K	12K–219K	~70K	Flash Attention in Triton, multimodal VLM from scratch, DeepSeek-R1, DPO / RLHF, Mamba/S4	🟢💤
Aladdin Persson	@AladdinPersson	92.3K	616–5.8K	~1.7K	Paper reviews, recommender foundation models, LLaMA4, career content	🟢💤
Jay Alammar	@arp_ai	64.6K	3.9K–222K	~42K	Transformer LLMs course, LLM agents w/ tool use, tokenizers, illustrated Word2Vec	🟢💤
CodeEmporium	@CodeEmporium	157K	511–11K	~2.8K	Transformers vs YOLO, CV timeline, diffusion, DALL-E, CLIP, ViT, DETR	🟢🔥
Venelin Valkov	@venelin_valkov	34.9K	❓ page crashed 3×	❓	Custom model benchmarking scripts, Python DL deployments, Hugging Face pipelines	🟢❓
mildlyoverfitted	@mildlyoverfi	8.05K	1.8K–29K	~8K	PyTorch paper-to-code (BentoML, RAG, NER)	🔴💤 (inactive 2+ yrs)
Abi Thakur	hakur	124K	430–182K	~25K	RAG / hybrid search / BM25 tutorials (4× Kaggle Grandmaster)	🟡🐢 (slowing)

Skip to content

Andrej Karpathy — The unanimous pick, and Reddit doesn't hedge here. From r/learnmachinelearning (u/Log_Dogg, +128 upvotes):

> *"Literally the goat, I can't think of a single better resource"*

And it's true: micrograd, makemore, "build the GPT tokenizer," reproducing GPT-2 — nothing else on the platform teaches you the transformer at that level of hands-in-the-guts intimacy. But look at the recency grade: 💤, last upload ~a year ago. That ~2.06M average is old evergreen gravity. Karpathy is a *curriculum*, not a subscription. Work through the back catalog like a book.

Sebastian Raschka — The *Build-an-LLM-from-Scratch* author, and the most alive of the from-scratch canon (🐢, not 💤). His architecture-comparison videos are the fastest way to understand what actually changed between model generations without reading six papers.

Umar Jamil — The deepest technical content in this tier, full stop. Flash Attention implemented in Triton. A multimodal VLM from scratch. DPO, RLHF, Mamba/S4. When people say "I want to *really* understand it," this is the channel I point them at — with the caveat that it's been ~a year since the last upload.

Jay Alammar — The illustrated-transformer guy. His visual explanations of tokenizers and Word2Vec are the on-ramp that makes Umar Jamil's content survivable. Dormant, evergreen, essential.

CodeEmporium — The surprise of this tier: still posting 🔥 weekly while the legends sleep. Diffusion, CLIP, ViT, DETR — solid mid-depth explainers with a computer-vision lean.

The rest, quickly: Aladdin Persson (92.3K) has drifted from PyTorch implementations toward paper reviews and career content. **Venelin Valkov** is Python-heavy and hands-on, but his page crashed all three scrapers — 34.9K subs confirmed, video metrics un[Skip to content]**overfitted** does beautiful paper-to-code work but has been ina**bhishek Thakur** (4× Kaggle Grandmaster) is slowing down but his RAG/BM25/hybrid-search tutorials remain practical.

Tier 3 — Applied ML, Bootcamp Courses & Model Tooling

End-to-end pipelines, roadmaps, framework courses. Where “I watched it” turns into “I built it.”

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
Krish Naik	@krishnaik06	1.46M	3.6K–496K	~63K	Agentic AI course, LangGraph / RAG, Claude Code, AgentOps, LLM guardrails/evals	🟡🔥
sentdex	@sentdex	1.44M	17K–308K	~64K	Frontier AI at home, open-source AI, Unitree G1 humanoid robotics, LLM agents	🟢🔥
CampusX	@campusx-official	638K	3.8K–68K	~25K	LLM evaluations series, Claude Code hooks/subagents, advanced RAG	🟢🔥
DeepLearning.AI	@Deeplearningai	681K	47K–935K	—	AI for Everyone, prompting course w/ Andrew Ng — the **active** Ng channel	🟢🔥
Alexander Amini	@AAmini	355K	8.7K–203K	~51K	MIT 6.S191 — CNNs, RNNs/transformers, deep generative, RL, AI for Science	🟢✅
Data School	@dataschool	261K	1.1K–15K	~3.5K	scikit-learn mastery, RAG, budget AI models	🟡💤
Daniel Bourke	@mrdbourke	250K	2.3K–240K	~40K	SLMs / on-device finetuning, DGX Spark, local multimodal RAG, “AI & ML Monthly”	🟢🐢
deeplizard	@deeplizard	169K	1.8K–14K	~6.1K	Stable Diffusion masterclass, computational graphs, AI art	🟡✅
Machine Learning w/ Phil	@MachineLearningwithPhil	45.1K	631–39K	~6.6K	Deep RL (PPO/DDPG/SAC/TD3), Ollama local LLMs, low-level programming	🟢💤
Jeremy Howard / fast.ai	@howardjeremyp	—	—	—	“Dangerous Illusion of AI Coding” interview, Answer.ai advocacy	🟢✅
Jeff Heaton	@JeffHeaton	96.2K	❓	❓	Legacy ML content, Second Life ML videos	🟡💤

⚠️ **A correction worth preserving: Andrew Ng’s *personal* channel (`@andrewyantakng`, 25.9K subs) is a dead Baidu-era archive** — the videos are 9–14 years old. People keep recommending it by name recognition. Don’t. The **DeepLearning.AI** org channel is where Ng actually publishes now, and it’s 🔥 active.

sentdex — Harrison Kinsley has been teaching Python-first ML for a decade, and ins he’s now running frontier AI at home and putting a Unitree G1 hu s paces. 1.44M subs, still 🔥, still 🟢. The rare mega-channel that aged into *more* interesting.

Skip to content

Alexander Amini — This is MIT 6.S191, an actual MIT deep learning course, free, updated, on YouTube: CNNs through transformers through RL through AI-for-Science. If you want a structured semester rather than a playlist mood-board, start here.

CampusX — The most current curriculum energy in this tier: an LLM evaluations series and *Claude Code hooks/subagents* content — course channels usually lag the frontier by a year; this one doesn't. (Hindi/English mix, deeply structured.)

Krish Naik — Enormous output across agentic AI, LangGraph, guardrails, evals. The 🟡 is for volume-over-editing — you'll filter, but there's real material here, and his cadence means he covers what shipped *this week*.

Daniel Bourke — On-device finetuning, DGX Spark experiments, local multimodal RAG, plus the "AI & ML Monthly" digest. Learn-in-public energy with genuinely useful engineering specifics.

Also worth knowing: Data School for scikit-learn fundamentals (💤 but evergreen), **deeplizard** for Stable Diffusion internals, **Machine Learning with Phil** for the deepest deep-RL back catalog on the platform (PPO/DDPG/SAC/TD3 implementations), and **Jeremy Howard** — the fast.ai philosophy is alive at Answer.ai, and his "Dangerous Illusion of AI Coding" interview is required contrarian listening for Part 5 people.

Tier 3.5 — Applied Data Science & Analyst Career

Career pathing, portfolio projects, analyst tooling. AI-adjacent, practically oriented.

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
Luke Barousse	@LukeBarousse	650K	1.2M–3.1M	~2M	Power BI & Excel full data-analyst courses	🟢✅
Nicholas Renotte	@NicholasRenotte	328K	7.2K–43K	~22K	RL Godot agent, LoRA finetuning, LangGraph trading agents	🟢🐢
Ken		278K	1.5K–44K	~11K	Pivoting toward AI-disruption / SaaS-building content	🟡✅
David Robinson	@data4democracy	15.9K	1K–5.1K	~2.4K	TidyTuesday / Riddler screencasts (dormant 6+ yrs)	🔴💤

Skip to content

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
Astroniz	@Astroniz	4.31K	65–3.4K	~500	NASA SPICE Python, asteroid/comet tracking, “AI in Astronomy”	🟢🐢

Look at **Luke Barousse’s** numbers again: 650K subs, **~2M average views**. His full-course format (Power BI, Excel) massively outperforms his subscriber base — the exact mirror image of the dormant legends. **Nicholas Renotte** is the most AI-forward here (RL agents in Godot, LoRA finetuning, LangGraph trading agents). **Ken Jee** is mid-pivot from pure data science toward AI-disruption content — grade accordingly. And **Astroniz** is the delightful micro-channel: NASA SPICE toolkit in Python, asteroid tracking, AI in astronomy — 4.31K subs of pure niche signal.

Part 2: The Frontier

Paper reviews, researcher interviews, safety, and the podcast circuit. How to stay research-aware without living on arXiv.

Tier 4 — Research Paper Reviews & Deep Technical Analysis

Primary-source walkthroughs. The “beyond-basics” tier, and the proof that depth can hold an audience.

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
Two Minute Papers	@TwoMinutePapers	1.83M	34K–320K	~110K	DeepSeek speed hacks, DeepMind/NVIDIA/Claude demos, fast paper awareness	🟡🔥
Yannic Kilcher	@YannicKilcher	326K	9.4K–172K	~45K	Paper analyses (TiDAR, Titans, GRPO/DeepSeekMath), “AGI is not coming!”, ML News	🟢🐢
MLST (Machine Learning Street Talk)	@MachineLearningStreetTalk	218K	4.9K–161K	~37K	Researcher interviews (Jumper, M. Jordan, Jeremy Howard), AGI/cognition	🟢🔥
bycloud	@bycloudAI	228K	18K–314K	~83K	DeepSeek architecture, LLM efficiency, JEPA, recursive LMs	🟢🔥
AI C	offeeBreak	64.3K	2.7K–45K	~11.5K	Flow-matching vs diffusion, energy-based transformers, decoding strategies	🟢💤
Discover AI	@code4AI	15.3K	1.1K–12K	~4.1K	Deep technical paper evaluations, model context limits, retrieval mechanics	🟢✅

Skip to content

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
hu-po	@hu-po	18.2K	1.3K–5.7K	~2.8K	Livestream paper deep-dives (RLHF, Gemini context, LDM)	🟡💤

bycloud — The standout of the tier, and my counter-example whenever someone claims technical depth can't scale on YouTube. DeepSeek's actual architecture, JEPA, recursive LMs, LLM efficiency — explained fast and funny without a single "SHOCKING" thumbnail, pulling an ~83K average. Technical-but-fun is a real lane, and he owns it.

Yannic Kilcher — The paper-review veteran. His walkthroughs (TiDAR, Titans, GRPO/DeepSeekMath) are where you go when the abstract wasn't enough but you're not ready to fight the appendix alone. Note the deliberately contrarian streak — "AGI is not coming!" — which I count as a feature: you need at least one skeptic in a feed full of exponential curves.

MLST — The most intellectually serious interview show in AI YouTube. Where else do you get John Jumper, Michael Jordan (the statistician, not the shooting guard), and Jeremy Howard talking cognition and AGI at length? Dense, philosophical, occasionally exhausting — in a good way.

Two Minute Papers — "What a time to be alive!" Károly's enthusiasm is a meme at this point, and the format has drifted toward demo-hype (hence 🟡), but as a *paper-awareness firehose* it still works: you'll know a result exists within days, then read it yourself.

Discover AI — The hidden gem here at 15.3K subs. Praised on Reddit for digging deeper into actual code and architecture parameters than mainstream channels bother to. Retrieval mechanics, real context-limit behavior — the stuff that decides whether your system works, covered at a depth the big channels skip.

Also: AI Coffee Break (Letitia) does lovely flow-matching-vs-diffusion explainers but has gone quiet (~8 months). **hu-po** livestreams multi-hour paper dissections — niche for

Skip to content

ormant lately.

Tier 4.5 — Visual / Technical ML Explainers

A compact bridge tier between implementation tutorials and paper-driven explanation.

Channel	Handle	Focus	Grade
Neural Breakdown with AVB	@avb_fj	LLMs, NLP, CV, RL, transformers, paper walkthroughs; latest = Unsloth DPO/SLM-alignment (~1 mo)	🟢✅
Gal Lahat	@GalLahat	Attention visualization, CNNs, LLM memory, hallucinations, infinite zoom, simulated systems	🟢🐢

Both crashed the metric scrape (subscriber/view figures unverified), both are worth your time anyway. **Neural Breakdown with AVB** is doing current-frontier walkthroughs (Unsloth DPO, SLM alignment) with animation quality that punches way above its size. **Gal Lahat** visualizes attention and LLM memory in ways that make hallucinations *make sense* mechanically.

Tier 5 — AI Safety, Alignment & Science Explainers

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
Computerphile	@Computerphile	2.63M	16K–583K	~108K	Shor's algorithm, why AI tokens are expensive, "Clever Hans" AI, post-quantum crypto	🟢🔥
Robert Miles AI Safety	@RobertMilesAI	170K	53K–360K	~189K	Alignment, mesa-optimizers, specification gaming	🟢✅

Small tier, zero filler. **Computerphile** remains the best "professor at a whiteboard with fanfold paper" channel in existence — the recent "why AI tokens are expensive" and "Clever Hans AI" videos are perfect send-to-a-colleague material. **Robert Miles** uploads rarely, but look at that ratio: 170K subs, ~189K average views. *His average video outperforms his entire subscriber count.* That's what happens when every upload is essential — mesa-optimizers and specification gaming explained so well that each vid[illegible] canonical reference.

Skip to content

Tier 6 — Long-Form Interview Podcasts

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
Lex Fridman	@lexfridman	5.02M	320K–1.7M	~879K	Broad interviews (Jensen Huang, physics, history); AI is a subset	🔴 ✅
Dwarkesh Patel	@DwarkeshPatel	1.36M	70K–904K	~309K	Deep technical AI interviews (training paradigms, Terence Tao + AI, chip design)	🟢 🔥
No Priors	@NoPriorsPodcast	88.2K	1.2K–46K	~15K	VC-adjacent founder/CEO interviews (Zuckerberg, Intel, OpenAI's Noam Brown)	🟡 🔥
TWIML AI Podcast	@twimlai	30.3K	—	—	Long-running AI/ML research interviews	🟢 ✅
Latent Space	@LatentSpaceTV	2.35K	46–3.7K	~423	"AI in Action" + "Paper Club" for AI engineers (audio-first)	🟢 ✅
Unsupervised Learning (Redpoint)	@RedpointAI	10K	400–12K	~2.5K	Jacob Effron interviewing AI founders on engineering constraints & scaling	🟢 ✅
The Robot Brains Podcast	—	—	—	—	Recommended in research passes; not verified as YT-first this pass	🟡 ❓

Dwarkesh Patel — The best pure-AI interview show running, period. He does the homework — actually reads the papers, actually understands the training-paradigm arguments — and the guests respond to that with answers you don't hear elsewhere. Terence Tao on AI, deep dives on chip design, ~309K average on genuinely technical conversations. This is my strongest "depth scales" data point.

Lex Fridman — Massive reach, and I've flagged him 🔴 anyway. AI is now a *subset* of an increasingly broad show (Jensen Huang next to history and physics episodes), the tone is polarizing, and technical follow-up questions are rare. There's value in the guest list; just know what you're getting.

Latent Space — Don't be fooled by the 2.35K YouTube subs — this is an audio-first podcast that happens to have a YT presence, and among actual working AI engineers it's near-canonical. The "Paper Club" and "AI in Action" formats are exactly the practitioner discourse the big shows can't do.

Also: No Priors for the VC/founder view (Zuckerberg, Noam Brown), **TWIML** as the lor[Skip to content]ch-interview institution, and **Unsupervised Learning** (Redpoint) for founders talking real engineering constraints.

Tier 6.5 — Business / Practitioner Podcasts & News Shows

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
Hard Fork	@hardfork	62.7K	4.7K–258K	~25K	Sundar Pichai / Satya Nadella interviews, OpenAI trial coverage	🟢🔥
AI For Humans	@AIForHumansShow	38.2K	5.7K–9.3K	~7.5K	GPT-5.6 Sol, Fable 5 survival, Claude/Alibaba spying story	🟢🔥
How I AI	@howiaipodcast	102K	3.8K–70K	~20K	Live harness-building, GPT-5.6 Sol benchmark, Claude Code loops	🟢🔥
The Artificial Intelligence Show	@aishowpod	—	—	—	Org adoption, agent security, business transformation, policy	🟢✅
Last Week in AI	@lastweekinai	5.89K	882–1.2K	~1K	Numbered weekly AI-news recap (#249, #248…)	🟢✅
This Week in AI	@ThisWeekinAIPodcast	8.2K	❓	❓	Alex Finn / Naveen Rao interviews (page crashed)	🟡❓
Practical AI	@practicalai_show	750	—	—	Podcast-first, minimal YT presence	🟡❓

The one to single out: **How I AI** — guests build their actual harnesses and Claude Code loops *live on screen*, which makes it the rare podcast that doubles as a Part 5 (agentic coding) resource. **Hard Fork** is your NYT-produced big-picture layer (Pichai, Nadella, the OpenAI trial), and **Last Week in AI** is the dependable numbered weekly recap for completionists.

Part 3: The News Diet

Model-watch and daily awareness — sorted by how much hype you'll have to filter. My rule: news channels are a routing layer to primary sources, never the verification layer.

Tier 7 — AI News, High Signal

Skip to content

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
AI Explained	@aiexplained-official	435K	41K–152K	~90K	Sourced model breakdowns (GPT-5.6, Claude Fable, Gemini), benchmarks, policy	🟢✅

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
Theo – t3.gg	@t3dotgg	549K	71K–170K	~117K	Dev-lens model reviews (GPT-5.6, Codex, local models, "moving to Linux")	🟢🔥

AI Explained — When a new frontier model drops, this is the first non-vendor analysis worth watching. From Reddit (u/astgabel, +33):

> *"the gold standard… down-to-earth and least hype-y channel for AI news"*

Every claim sourced, benchmarks actually contextualized, policy coverage that isn't doomer-bait. One standing caveat from the same threads: a commenter argues the channel has "changed to a sales funnel for paid content" (there's a paid tier now). I still rate the free videos as the best model-release analysis on the platform — but you deserve the whole picture.

Theo (t3.gg) — The developer's lens on model news. Where AI Explained asks "what does this benchmark mean," Theo asks "what happens when I point this at a real codebase." Opinionated, fast, occasionally wrong out loud and corrected in public — which I trust more than never-wrong-sounding.

Tier 7.5 — AI News, Daily / Format-Separated

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
AI Daily Brief	@AIDailyBrief	584K	3.3K–11K	~6K	Daily AI-economy news, model-release recaps (Nathaniel Whittemore)	🟢🔥
Nate B Jones	@NateBJones	305K	16K–90K	~35K	Near-daily AI strategy, agent economics, enterprise adoption, Claude memory build	🟢🔥
Asianometry	@asianometry	944K	~202K recent; historic 50K–1.2M	~180K	Semiconductor / AI-hardware history, chip testing, ASML/TSMC, boom-bust cycles	🟢✅
What's AI (Louis Bouchard)	@WhatsAI	73.1K	259–9.3K	~2K	Loop / harness-engineering explainers, AI-engineering foundations course	🟢✅
Dr Alan D. Thompson	@DrAlanDThompson	59.6K	1K–26K	~6K	"First look" humanoid robots (Figure 03, Xiaomi CyberOne), ASI tracking	🟢✅

Skip to content

Nate B Jones deserves a spotlight, partly because my own pipeline nearly buried him: one research pass mis-scraped him at 6K subs. The real figure is **305K**, near-daily uploads, and he occupies a lane almost nobody else does — AI *strategy*: agent economics, enterprise adoption, what these releases mean for how organizations actually work. The recommended complement to the whole Tier 11 automation crowd, minus the "$20K/month" framing.

AI Daily Brief is the lowest-effort way to maintain daily awareness (Nathaniel Whittemore, calm, ~15 minutes). **Asianometry** is cross-listed here but lives in Tier 16 — more below. **What's AI** and **Dr Alan D. Thompson** round out the awareness layer (the latter is your humanoid-robot release tracker).

Tier 8 — AI News & Tool Roundups, Mainstream (hype-adjacent)

Useful for day-1 awareness. Also the tier where Reddit's receipts come out, so let's be precise about who's flagged for what.

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
Fireship	@Fireship	4.23M	292K–1.0M	~663K	Rapid dev/AI news, "code report", OSS incidents, systems concepts	🟡🔥
IBM Technology	@IBMTechnology	1.74M	3.8K–85K	~19K	Agentic frameworks, MCP vs Skills, KV cache, AI security explainers	🟢🔥
Matt Wolfe	@mreflow	978K	38K–109K	~82K	Weekly "AI News" + tool roundups (FutureTools.io)	🟡🔥
AI Search	@theAIsearch	704K	73K–498K	~166K	Weekly AI news + tool/model demos, image/video gen	🟡🔥
Matthew Berman	@matthew_berman	623K	42K–183K	~98K	Model testing/news	🔴🔥
TheAIGrid	@TheAiGrid	396K	1.9K–77K	~17K	Daily model leaks/news	🔴🔥
Wes Roth	@WesRoth	322K	19K–155K	~64K	Paper/news coverage	🔴🔥
MattVidPro AI	@MattVidPro	300K	4.6K–44K	~13K	Model/tool testing (GPT-5.6, Fable 5, ElevenLabs, Krea)	🟡✅
David Shapiro	@DaveShap	189K	7.6K–36K	~17.5K	Post-labor economics, UBI, AGI timelines	🔴🔥
1littlecoder	@1littlecoder	110K	929–12K	~3.6K	Fast, practical model tutorials (Claude, OCR, GPT, Nemotron)	🟡✅

Skip to content

Ab Matt Wolfe + Matthew Berman + Wes Roth get recommended together constantly — they're the algorithmic starter pack for anyone who watches

one AI video. The Reddit record on each, verbatim:

On **Matthew Berman** (u/fasti-au):

> *"He can't code so what he sees is repeated"*

On **Wes Roth** (u/zackler6, +86 — note this one's a *defense*):

> *"videos are actually good… but titles over the top"*

On **TheAIGrid** (u/MysteriousPepper8908, +25):

> *"low-quality content mill"*

On **Matt Wolfe** — the most defended of the group (u/Substantial-Comb-148, +8):

> *"my go-to for quick weekly rundowns… always transparent about affiliations upfront"*

And on **David Shapiro** (u/laudanus, +16):

> *"eccentric hobbyist… no academic ML background"*

My honest take: these channels solve a real problem (velocity — knowing something dropped within hours) while creating another (framing — everything is SHOCKING, INSANE, or HERE (WOAH)). Use them as *notification systems*, then verify with Tier 7 or the vendor channels. Never let them be your only layer.

Two exceptions in this tier worth separating from the pack: **Fireship** is hype-*paced* but not hype-*brained* — the "code report" format compresses real information density into 4 minutes, and 663K average views say the format works. **IBM Technology** is the anti-hyp[Skip to content]borate channel producing genuinely solid whiteboard explainers on agentic frameworks, MCP vs Skills, and KV caching, at 🔥 cadence.

Part 4: The Hardware Truth

Local LLMs, runtime benchmarking, VRAM reality, and the silicon underneath. The r/LocalLLaMA core.

Tier 9 — Local LLMs, Runtime & Hardware Benchmarking

Real benchmarks, tok/s, quantization, local rigs — where marketing claims go to be measured.

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
Fahd Mirza	@fahdmirza	763K	357–7.6K	~2.7K	High-cadence local model tests, GPU comparisons ($2000 Huawei GPU vs NVIDIA)	🟢🔥
Alex Ziskind	@AZisk	531K	58K–324K	~117K	Apple Silicon vs DGX Spark/AMD, local AI OS impact, 4-bit, tok/s	🟢🔥
Donato Capitella	@donatocapitella	99.7K	❓ (video list crashed)	❓	Strix Halo, Radeon AI Pro 9700, llama.cpp, vLLM, ROCm, agentic-AI security	🟢🔥 (last upload ~5d)
Digital Spaceport	@DigitalSpaceport	93.9K	2.9K–81K	~31.5K	Local AI server builds, motherboard/CPU combos, Gemma/Qwen benchmarks	🟢🔥
Julia Turc	@juliaturc1	72.1K	25.5K–36.3K	~30K	*(cross-listed from Tier 1)* efficient serving, FP4 quant, MoE	🟢✅
Bijan Bowen	@bijanbowen	64.8K	16K–48K	~35K	"First Test / Hands-On" open + local coding model reviews — contested depth	🟡🔥
Mukul Tripathi	@MukulTripathi	—	—	—	RTX Pro 6000 / Blackwell, vLLM, high-context inference, home AI servers	🟢🐢
Token Chaser	@tokenchaser	7.5K	3.4K–35K	~9K	Rapid local-vs-cloud head-to-head battles (Qwen3.6, Fable 5, Opus), API-endpoint deploys	🟢🔥
Protorikis	@Protorikis	12K	900–14K	~4.2K	Quantization perplexity tests, 16→4→2-bit degradation, GGUF/EXL2	🟢
Codacus	@Codacus	11K	800–11K	~3.5K	Running models on older consumer GPUs, modest-hardware optimization	🟢
Luke's Dev Lab	@lukesdevlab	6K	400–5.2K	~1.6K	Custom model API servers, local context loading, inference optimization	🟢
Tonbi's AI Garage	@TonbisAIGarage	5K	300–4.1K	~1.3K	SD parameters, local WebUI configs, custom finetunes	🟢
Lev…	@Level1Tech	❓	—	—	Local AI server hardware, PCIe/GPU virtualization (crashed every scrape pass)	🟡❓

Skip to content

The Benchmark Caveat: Read Before Buying a GPU

Before you take *any* number from *any* of these channels to the checkout page, here's the emblem of this whole tier. In the comments of Donato Capitella's Radeon AI Pro 9700 video, a Redditor reported **~40 tok/s** for `gpt-oss:20b` at 4-bit — where the video showed **130 tok/s**. Same model, same card, **3× apart.** The likely culprit: runtime. Ollama vs llama.cpp vs vLLM can legitimately differ by that much on identical hardware.

Treat every hardware benchmark as a configuration-specific data point, not a universal result. The channel isn't lying; your stack just isn't their stack.

Alex Ziskind — The biggest genuinely rigorous channel in the tier. Apple Silicon vs DGX Spark vs AMD, measured in actual tok/s with methodology on screen. If you're deciding between a Mac Studio and a dedicated GPU box, his back catalog *is* the buyer's guide.

Donato Capitella — The reconciliation story here matters: an early pass had him stale at 8K subs; the truth is **99.7K** and posting within the last 5 days. Strix Halo, ROCm, llama.cpp vs vLLM — plus a genuinely distinctive second lane in *agentic-AI security*, which almost nobody else covers with engineering rigor.

Token Chaser — 7.5K subs and one of my favorite follows in the tier. The format: rapid local-vs-cloud head-to-heads (Qwen3.6 vs Fable 5 vs Opus) run as *multi-step qualitative tests* — actual coding tasks, actual HTTP API deployments — rather than the one-shot "write me a snake game" prompts everyone else calls a benchmark. (Alias note from the dedup pass: "Token Chaser" and "Token Chasers" are the same creator.)

Protorikis — The quantization specialist. Perplexity degradation from 16-bit → 4-bit → 2-bit, GGUF vs EXL2, measured properly. When you're wondering whether Q2 is actually usable for your use case, this is where the real answer lives.

Digital Spaceport — Full local-AI *server* builds: motherboard/CPU combos, multi-GPU layouts, then Gemma/Qwen benchmarks on the finished rig. The homelab end of the spectrum.

Skip to content

And the cautionary tale: Fahd Mirza — 763K subscribers, ~2.7K average views. That's the most extreme badge-vs-reach gap in this entire directory. To be clear, the 🟢

stands: high-cadence, hands-on, honest testing (the $2000 Huawei GPU vs NVIDIA comparison is exactly the content nobody else makes). But his audience clearly treats him as a search index, not a feed — and honestly, that might be the right way to use him.

Smaller but real: Codacus (models on older consumer GPUs — the "average Joe hardware" niche), **Luke's Dev Lab** (custom model API servers), **Tonbi's AI Garage** (local WebUI configs and finetunes), **Mukul Tripathi** (RTX Pro 6000 / Blackwell-class home servers, high-context inference). **Bijan Bowen** posts fast "first test" reviews of every open model drop — useful velocity, contested depth (🟡). **Level1Techs** crashed every scrape pass but the forum-adjacent hardware content is well-regarded ❓.

Tier 16 — AI Hardware & Semiconductor History

The physical substrate — chip design, fabs, lithography, boom/bust cycles.

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
Asianometry	@Asianometry	944K	50K–1.2M	~180K	Semiconductor fab history, LISP machines, "why Soviet computers failed", ASML/TSMC, AI boom-bust	🟢✅

A one-channel tier, because **Asianometry** has no peer. Outstandingly researched essays on the hardware economics powering AI: ASML's monopoly, TSMC's rise, LISP machines, why Soviet computing failed, and — most relevant right now — the anatomy of semiconductor boom-bust cycles. If you want to understand whether the current AI capex wave is 1999 or 2004, this is the channel doing the archival work. Cross-listed in the news tiers for relevance, but this is its true home.

Part 5: The Builders

Skip to content

Ag…ness engineering, and the business of automation. By posting cadence, this is the hottest zone in AI YouTube right now — nearly everyone here is 🔥,

and they're all converging on the same frontier: Claude Code, Codex, and multi-agent orchestration.

Tier 10 — Agentic Coding: Claude Code / Codex / Harness Engineering

Where real software development meets agent loops, context engineering, and multi-agent orchestration.

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
Cole Medin	@ColeMedin	216K	1.7K–131K	~32K	Harness engineering, agent orchestration, Pydantic AI 2.0, Karpathy LLM Wiki	🟢🔥
Matt Pocock	@mattpocockuk	295K	5.2K–85K	~32K	Building "shared domain language" w/ agents, testable architectures, deep modules	🟢
AI Jason	@AIJasonZ	227K	6.8K–206K	~47K	Coding agents, Loop/Harness engineering, MCP, "Anthropic killed tool calling"	🟢✅
IndyDevDan	@indydevdan	136K	16K–157K	~34K	Multi-agent orchestration, Claude Code/Pi Agent, agentic security, local MLX stack	🟢🔥
Sam Witteveen	@samwitteveenai	124K	7.8K–82K	~28K	Model breakdowns + agent/LLM engineering w/ Colab notebooks	🟢🔥
Brian Casel	@briancasel	69.5K	2.1K–753K	~40K	Claude Code CRM / time-tracker builds, git worktrees for parallel agents	🟢🔥
GosuCoder	@GosuCoder	27.1K	4.1K–61K	~15K	Weekly AI coding-agent shootouts (GLM, Gemini 3, Opus, Cursor)	🟢🔥
BMad Code	@BMadCode	34.2K	❓	❓	BMad-Method agentic-dev framework tutorials (page crashed)	🟢❓
Armin Ronacher	@ArminRonacher	9.15K	4.1K–13K	~8K	"State of Agentic Coding" monthly series (Flask creator)	🟢🐢
Matt Maher	@MetalSole	3K	200–5K	~1.1K	Python scripts, local AI coding basics, beginner-friendly environments	🟢

This tier is personal for me — it's my daily workflow — so let me be specific about who earns which slot.

IndyDevDan — The repeated Reddit pick for "best Claude Code content outside Anthropic's own," and I co-sign. Multi-agent orchestration, agentic security, a local MLX stack — he's consistently 2–3 months ahead of where the median tutorial channel is, and he ships the workflows he preaches.

Skip to content

Co — The systematizer. Where others demo, Cole builds *frameworks* for thinking about harness engineering and agent orchestration (his Pydantic AI 2.0

material is the reference). Posting within 6 hours of my scrape — the cadence matches the frontier.

Sam Witteveen — From Reddit (u/hassan789_, +17):

> *"zero fluff, zero BS"*

That's the whole review. Model breakdowns and agent engineering, every video with a Colab notebook you can run. The highest signal-to-runtime ratio in the tier.

GosuCoder — 27.1K subs, weekly coding-agent shootouts (GLM vs Gemini 3 vs Opus vs Cursor), and — more importantly — an explicit methodology: structured context management, spec files, and code review over vibes. His "don't vibe to production" stance is the necessary adult supervision for this entire content zone.

Matt Pocock — The renowned TypeScript educator pivoted to how AI reshapes software engineering itself: building a "shared domain language" with agents, testable architectures, deep modules. The most *software-engineering-brained* channel in the tier — agents as a design problem, not a magic trick.

Brian Casel — Real products built on camera (a CRM, a time tracker), and the single most useful workflow demo I've seen this year: **git worktrees to keep multiple Claude Code agents active in parallel** on separate branches, reviewed and merged independently. One video of his hit 753K views — the appetite for real builds is enormous.

Armin Ronacher — Yes, *that* Armin Ronacher (Flask, Jinja2). His monthly "State of Agentic Coding" series is what it looks like when someone with two decades of tooling credibility evaluates this space: measured, skeptical, zero affiliate links. 9.15K subs is criminal underexposure.

Also: AI Jason for MCP and provocations that are better-researched than their titles ("A[Skip to content]l calling"), **BMad Code** for the BMad-Method framework (page cra**laher** for gentle beginner on-ramps.

Tier 10.5 — AI-Engineer Conference & Practitioner Talks

Channel	Handle	Latest-upload themes	Grade
AI Engineer	@aiDotEngineer	Agent skills, AI coding workflows, local AI, context engineering, complex codebases	🟢🔥

AI Engineer is the talk archive of the AI Engineer conferences — speakers from Anthropic, OpenAI, NVIDIA, MIT, uploads landing within hours-to-days of the events. It's the best complement to the whole tier above: where individual creators show you *their* workflow, this shows you fifty practitioners' workflows in parallel. Also the single best channel for discovering people who belong in the next edition of this directory.

Tier 11 — AI Agents Business, Automation & n8n

Build / monetize / scale agent pipelines. Fair warning on framing: this tier skews hard toward outcomes and "$/month" packaging — the engineering is often real, but the business claims deserve your skepticism.

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
Tina Huang	@TinaHuang1	1.25M	30K–753K	~220K	NotebookLM 2.0, Claude Cowork setup, local AI agents in 26 min	🟢🔥
Nate Herk	@nateherk	853K	8.4K–175K	~78K	n8n + Claude Code automation, agent loops, monetization	🟡🔥
Liam Ottley	@LiamOttley	819K	8.3K–211K	~55K	AI-automation-agency business model, "$20K/mo with Claude Code"	🟡🔥
Greg Isenberg	@GregIsenberg	670K	31K–157K	~83K	"AI agents are the new SaaS", startup ideas, solo-agent business	🟡🔥
Nick Saraev	@nicksaraev	469K	9.7K–210K	~81K	Agentic workflows, Claude Code apps	🔴🔥
Sabrina Ramonov	@sabrina_ramonov	318K	1.8K–152K	~20K	AI-money-making, Claude+Canva workflow, near-daily uploads	🟡🔥
Dave Ebbelaar	@daveebbelaar	275K	3.7K–61K	~18K	Python for agents, agentic RAG from scratch, DS→AI engineer, Pydantic	🟢🔥
Riley Brown	@rileybrownai	259K	8.3K–88K	~40K	Vibe coding, Codex/Claude/Cursor, AI assistant builds	🟡🔥
All About AI	@AllAboutAI	225K	1.4K–10K	~4.7K	Agentic AI trading (Polymarket/Hyperliquid), MCP, automation	🟡🔥
Me Pra		81.5K	495–38K	~6.7K	PraisonAI, local agents w/ Ollama, Claude Code + Slack	🟡✅

Skip to content

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
James Briggs	@jamesbriggs	81.4K	1.8K–10K	~4.3K	LangChain / OpenAI Agents SDK, RAG, vector DBs	🟢💤 (~9 mo)
Raj Amjad	@RAmjad	12K	800–8.1K	~3K	Opinionated agentic frameworks, beginner enterprise automations	🟡

The two 🟢 grades in a tier full of 🟡: **Tina Huang** earns hers by being the rare mega-channel (1.25M, ~220K avg) that teaches workflows without inflating outcomes — her "local AI agents in 26 minutes" delivers what the title says. **Dave Ebbelaar** earns his by staying an engineer in a marketer's tier: Python-first, agentic RAG built from scratch, Pydantic patterns, honest about the data-scientist-to-AI-engineer path.

One 🔴 to explain: **Nick Saraev**'s content itself is competent, but Reddit threads flag his Maker Skool as over-marketed — the classic funnel pattern where the free content is the ad. Watch accordingly.

And a routing note: if what you actually want from this tier is "how will agents change *work*" rather than "how do I sell agents," that's **Nate B Jones** back in Tier 7.5 — enterprise-workflow analysis without the income-claim framing.

Tier 11.5 — Non-Technical / End-User AI Adoption

Tool fluency and knowledge-work workflows. Not a rigorous ML or agent-engineering core — this is the tier you send to your colleagues, not the one you watch yourself.

Channel	Handle	Focus	Grade
The AI Advantage	@aiadvantage	ChatGPT/Claude/Midjourney tutorials, prompting, AI workflows	🟡 (posts within ~1 yr)
Natalie Lambert / GenEdge	@NatalieLambert-GenEdge	AI marketing workflows, NotebookLM, AI content teams, copyediting	🟡💤 (uploads 1–2 yrs old)

The AI Advantage is the standard answer to "my non-technical teammate wants to get good at ChatGPT/Claude — where do they start?" **Natalie Lambert** covers the marketing-team angle (NotebookLM workflows, AI content operations), though her up

Skip to content

stalled.

Part 6: The Creators

Image gen, AI film, creative coding, robotics, and the vendors themselves.

Tier 12 — ComfyUI / Stable Diffusion / Image Generation

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
Olivio Sarikas	@OlivioSarikas	250K	1.4K–15K	~7.4K	Seedance 2.0, Midjourney, Nano Banana, gen workflows	🟡🐢
Sebastian Kamph	@sebastiankamph	182K	1.7K–23K	~7.5K	ComfyUI course + node guides, model shootouts (Z-Image, Kling)	🟢✅
Nerdy Rodent	@NerdyRodent	64.8K	3.1K–43K	~11.3K	ComfyUI workflows (Krea-2, LTX, TTS, music) — scripted, non-hype	🟢🔥
Scott Detweiler	@sedetweiler	59.8K	1.5K–149K	~50K	ComfyUI/SDXL/ControlNet/LoRA (Stability.ai PM)	🟢💤
SECourses (Dr. Furkan)	@SECourses	53K	❓	❓	FLUX/SDXL full finetuning & DreamBooth master tutorials (page crashed)	🟢❓
Latent Vision	@latentvision	38.5K	13K–61K	~30K	ComfyUI / IPAdapter deep technical dives	🟢✅
Benji's AI Playground	@BenjisAIPlayground	32.5K	❓	❓	ComfyUI tutorials; personally polarizing per Reddit (page crashed)	🟡❓

The correction I flagged in the intro lands here: **Latent Vision** — unverified in my first pass, now confirmed at 38.5K subs with a ~30K average. That near-1:1 ratio makes sense once you know who it is: Matteo, the *developer of IPAdapter itself*, explaining ComfyUI internals. This isn't a tutorial channel that uses the tools; it's the toolmaker teaching. The classic Reddit "SD starter pack" cites four names together — **Sebastian Kamph** + **Olivio Sarikas** + **Latent Vision** + **Nerdy Rodent** — and that's still the right four, with one trend-line caveat: this whole zone is quietly cooling at the top. Olivio has slowed to 🐢, Scott Detweiler (a Stability.ai PM, and the best ControlNet/LoRA back catalog around) has gone 💤. The active reliable core is now Kamph + Nerdy Rodent + Latent Vision.

Tier 13 — AI Film & Video Generation

Skip to content

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
Mickmumpitz	@mickmumpitz	181K	39K–600K	~200K	Free/local AI film pipelines, consistent characters, ComfyUI + Blender VFX	🟢✅
Theoretically Media	@TheoreticallyMedia	191K	14K–47K	~28K	AI video tools (Seedance, Kling, Runway, Google Omni), production tests	🟡🔥

Mickmumpitz is the one to watch here — free/local AI film pipelines with *consistent characters* (the actual hard problem), ComfyUI wired into Blender VFX, ~200K average views. **Theoretically Media** is your tool-coverage layer: Seedance, Kling, Runway, tested against production use rather than cherry-picked demos.

Tier 14 — Official Vendor Channels (source of truth)

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
OpenAI	@OpenAI	1.99M	3.1K–125K	~23K	"ChatGPT Work" enterprise suite, Codex, GPT-5.6 use cases	🟢✅
Google DeepMind	@GoogleDeepMind	902K	7.4K–588K	~179K	Gemini 3 Deep Think, science AI (drug-resistance, WeatherNext), agents	🟢✅
Anthropic	@anthropic-ai	732K	6.2K–806K	~315K	Claude Fable 5 / Opus, interpretability, Cowork, MCP, safety	🟢✅
Hugging Face	@HuggingFace	137K	1.3K–78K	~14.4K	Coding agents (Tau), LeRobot, MoE, RoPE, ML Club/Podcast	🟢✅

Subscribe to all four — it's marketing, but it's *primary-source* marketing, and when a model drops the vendor video is ground truth for what was actually claimed. The engagement numbers tell their own story: **Anthropic averages ~315K views** (interpretability research and safety content that people genuinely watch), **DeepMind ~179K** (the science-AI work — drug resistance, WeatherNext — is legitimately great), while **OpenAI averages just ~23K on 1.99M subs** — heavily enterprise-oriented output that even their own audience skips. **Hugging Face** is the sleeper: ML Club sessions and RoPE/MoE explainers that function as a free grad seminar.

Tier 15 — Creative AI & "Weird Programs"

Skip to content

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
Emergent Garden	@Emergent Garden	273K	14K–1.7M	~250K	“Recursive Self-Improvement”, AI-in-Minecraft (Mindcraft), cellular automata, “I Quit My Job to Make Weird Programs”	🟢✅

Another one-channel tier that refuses to be merged into anything else. **Emergent Garden** makes narrative-driven programming videos about emergence and complexity — LLM agents surviving in Minecraft (the Mindcraft project), cellular automata, recursive self-improvement — and the flagship “I Quit My Job to Make Weird Programs” tells you everything about the vibe. Extremely high signal wrapped in genuine playfulness. 273K subs, ~250K average: near-perfect engagement.

Tier 17 — Robotics & Maker Engineering

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
Michael Reeves	@MichaelReeves	7.8M	5.2M–13M	~8M	Dog catapult, scam-bot, goldfish stock-trading (comedic robotics)	🟢💤
Stuff Made Here	@StuffMadeHere	4.75M	2.2M–7.3M	~5M	Robot golf, autocorrect putter, high-production engineering builds	🟢🐢
Code Bullet	@CodeBullet	3.46M	1.3M–3.5M	~2.4M	AI plays games, “same game” dev challenges	🟡🐢
Sebastian Lague	@SebastianLague	1.4M	142K–6.1M	~800K	Coding neural nets from scratch, Rubik’s solver, fluid/ray-tracing sims	🟢🔥
James Bruton	@jamesbruton	1.4M	47K–256K	~100K	Ball-balancing robots, 5-servo biped, ESP32 hexapod — open code/CAD	🟢🔥
DroneBot Workshop	@Dronebotworkshop	677K	5.3K–150K	~50K	ESP32 TTS/OTA/low-power, Arduino Uno Q, LiDAR sensors	🟢🔥
Paul McWhorter	@paulmcwhorter	440K	883–1.6K	~1.2K	“AI on the Edge” series — OpenCV facial recognition, object tracking	🟢🔥
Skyentific	@Skyentific	116K	8.7K–28K	~17K	EtherCAT robot comms, NVIDIA Isaac Lab bipedal sim-to-real	🟢🐢
Articulated Robotics	@ArticulatedRobotics	75.5K	❓	❓	ROS tutorials, mobile-robot build series (page crashed)	🟢❓

Robotics YouTube skews entertainment-heavy at the top — Reeves at an 8M average, Stuff Made Here at 5M, Code Bullet at 2.4M — glorious to watch, thin to learn from. For *rig* **kyentific** (EtherCAT comms, NVIDIA Isaac Lab sim-to-real), **Articulated Robotics** (the ROS tutorial series everyone recommends), and **Paul**

Skip to content

McWhorter's "AI on the Edge" (OpenCV facial recognition on real hardware, taught with infinite patience at a ~1.2K average that deserves 100×).

The special case: **Sebastian Lague**. His "Coding Adventures" — neural nets from scratch, fluid sims, ray tracing — are the most beautiful programming videos ever made. Full stop. And one honest flag on **James Bruton**: universally praised for open-sourcing every build's code and CAD, but several Redditors note the *explanations* stay surface-level past Arduino basics. Watch for inspiration, read his repos for depth.

Part 7: The Graveyard and the Blacklist

Who to stop expecting new content from, who never earned a slot, and the one name that's an active warning.

Tier 18 — Research Interviews & Foundations (largely legacy/sparse)

Channel	Handle	Subs	Recent views (min–max)	Avg	Latest-upload themes	Grade
Aleksa Gordić — The AI Epiphany	@TheAIEpiphany	64.5K	1.6K–31K	~5K	Interviews w/ Groq, HuggingFace, Meta, DeepMind researchers	🟡💤
Connor Shorten	@connor-shorten	52.3K	1K–82K	~15K	DSPy, Weaviate, RAG explainers	🔴💤 (1+ yr)
Michael Bronstein	@MichaelBronsteinGDL	14.4K	❓	❓	Geometric Deep Learning course lectures	🟡🐢
Arxiv Insights	@ArxivInsights	103K	~50.9K historic	—	The famous "why humans learn faster than AI" — Reddit literally asks "what happened?"	🔴💤
Henry AI Lab	@hdnh2610	1.6K	❓	❓	Curated DL-paper video archive; tiny/legacy	🔴💤
Brandon Rohrer	@BrandonRohrer	89.8K	300–9.6K	~1.5K	"How Data Science Works", CNN-from-scratch series	🔴💤

Worth one line each: **Aleksa Gordić** recorded genuinely good researcher interviews (Groq, DeepMind, Meta) before going quiet. **Michael Bronstein's** geometric deep learning lectures remain the reference for that subfield. **Brandon Rohrer's** CNN-from-scratch teaches convolutions better than most current content. Legacy ≠ worthless — it just means don't wait for episode 2.

Skip to content

The Dormant and Dead List

Decent sub counts, confirmed inactive as *current* resources. Do NOT treat these as live feeds:

Channel	Subs	Last activity	Status
Arxiv Insights	103K	years	🔴💤 legacy conceptual gold, no new uploads
Connor Shorten	52.3K	1+ yr	🔴💤
Henry AI Lab	1.6K	legacy	🔴💤
Brandon Rohrer	89.8K	dormant	🔴💤
mildlyoverfitted	8.05K	2+ yrs	🔴💤
David Robinson	15.9K	6+ yrs	🔴💤
Andrew Ng (personal)	25.9K	9–14 yrs	🔴💤 use DeepLearning.AI instead
Karpathy	1.57M	~1 yr	🟢💤 reference library, not a feed
Umar Jamil / Jay Alammar / Aladdin Persson	85K / 64.6K / 92.3K	~1 yr	🟢💤 evergreen but inactive
AI Coffee Break	64.3K	~8 mo	🟢💤
James Briggs	81.4K	~9 mo	🟢💤
Brandon Foltz / zedstatistics / Reducible	—	1–3 yrs	🟡💤 prerequisite supplements only

Note the two different 💤 flavors: the 🔴💤 rows are channels whose content has aged past usefulness or whose communities have moved on. The 🟢💤 rows — Karpathy above all — are *libraries*: dormant as feeds, permanent as curricula.

The Blacklist

One name, and it's non-negotiable: **Siraj Raval**. Repeatedly flagged for plagiarism (including of an academic paper) and course-refund fraud. From Reddit (u/new_name_who_dis_, +18):

> *"I remember when this sub used to hate this guy for all his fraud"*

The complicated truth: his early videos genuinely motivated a lot of beginners into ML, an[Skip to content]eople are researchers now. But motivation isn't trust, and there are ter 1–2 of this guide that teach the same material without the plagiarism. Not a trusted source.

Candidates Discovered but Not Admitted

For completeness — names that surfaced in Reddit recommendations but failed my verification bar. Listing them so you know they weren't *missed*, they were *checked*:

Candidate	Reason not admitted
@cognibuild / Cognibuild AI	YouTube page repeatedly failed to render for content/activity verification
@spatialwebai	One Reddit link only; page didn't render enough to establish focus/identity
@g0t4	Resolves more to a developer/GitHub presence than a validated AI channel
SwissCognitive	No canonical YT page recovered
Varun Mayya	Broader career/business creator, not clearly AI-first
MarkTechPost	Primarily a publication/site signal, not a validated YT channel
DailyDoseofDS	No canonical YT channel / reliable current context recovered

Part 8: Subscription Strategy

You don't need 185 channels. You need the right 8–12 for what you're actually doing. Pick your stack.

The Curated Stacks

🛠️ **Agentic engineering (if you ship code with agents daily):** Cole Medin · IndyDevDan · Sam Witteveen · AI Jason · Matt Pocock · GosuCoder · Brian Casel · AI Engineer (conf) · Anthropic · OpenAI

🖥️ **Local models & hardware truth (before you spend money on GPUs):** Donato Capitella · Alex Ziskind · Fahd Mirza · Mukul Tripathi · Digital Spaceport · Token Chaser · Protorikis · Julia Turc · bycloud · Level1Techs

🔬 **Research awareness without living on arXiv:** AI Explained · bycloud · Yannic Kilcher · MLST · Two Minute Papers · Dwarkesh Patel · Discover AI · Neural Breakdown w/ AVB · Google DeepMind

Skip to content

🧮 **Foundations from zero (the self-taught degree):** 3Blue1Brown · StatQuest · Karpathy · Sebastian Raschka · Umar Jamil · Serrano Academy · Welch Labs · Steve

Brunton · Alexander Amini · DeepLearning.AI

🎨 **Image/video systems:** Latent Vision · Sebastian Kamph · Nerdy Rodent · Olivio Sarikas · Mickmumpitz · Theoretically Media · SECourses

📰 **Daily awareness (routing only — verify elsewhere):** AI Daily Brief · Nate B Jones · AI Explained · Hard Fork

Master Leaderboard: Top 40 by Subscriber Count

For the record — and as a standing exhibit of why this metric alone tells you almost nothing (compare #25 Fahd Mirza's 2.7K average against #26 Anthropic's 315K):

#	Channel	Subs	Avg recent views	Tier
1	3Blue1Brown	8.47M	1.69M	T1
2	Michael Reeves	7.8M	8M	T17
3	Lex Fridman	5.02M	879K	T6
4	Stuff Made Here	4.75M	5M	T17
5	Fireship	4.23M	663K	T8
6	Code Bullet	3.46M	2.4M	T17
7	Computerphile	2.63M	108K	T5
8	OpenAI	1.99M	23K	T14
9	Two Minute Papers	1.83M	110K	T4
10	IBM Technology	1.74M	19K	T8
11	StatQuest	1.66M	~6K*	T1
12	Andrej Karpathy	1.57M	2.06M	T2
13	Krish Naik	1.46M	63K	T3
14	sentdex	1.44M	64K	T3
15	James Bruton	1.4M	100K	T17
16	Sebastian Lague	1.4M	800K	T17
17	Dwarkesh Patel	1.36M	309K	T6
18	Tina Huang	1.25M	220K	T11
19	Matt Wolfe	978K	82K	T8
		944K	180K	T16
		902K	179K	T14
22	Welch Labs	887K	~130K	T1

Skip to content

#	Channel	Subs	Avg recent views	Tier
23	Nate Herk	853K	78K	T11
24	Liam Ottley	819K	55K	T11
25	Fahd Mirza	763K	2.7K	T9
26	Anthropic	732K	315K	T14
27	AI Search	704K	166K	T8
28	DeepLearning.AI	681K	—	T3
29	DroneBot Workshop	677K	50K	T17
30	Greg Isenberg	670K	83K	T11
31	Luke Barousse	650K	2M	T3.5
32	CampusX	638K	25K	T3
33	Matthew Berman	623K	98K	T8
34	AI Daily Brief	584K	6K	T7.5
35	Theo (t3.gg)	549K	117K	T7
36	Steve Brunton	540K	14K	T1
37	Alex Ziskind	531K	117K	T9
38	Nick Saraev	469K	81K	T11
39	Paul McWhorter	440K	1.2K	T17
40	AI Explained	435K	90K	T7

(Ranks 41–185 are fully tabulated within their tier sections above; StatQuest average is partial — only 3 recent view counts rendered.)

Cross-Cutting Insights

If you skimmed everything above, these seven findings are the takeaway:

1. **Subscribers ≠ current reach.** The sharpest example bears repeating: Fahd Mirza (763K subs, ~2.7K avg) vs Anthropic (732K subs, ~315K avg) — **~115×** apart at the same badge. Judge live relevance by recent-view average × recency, never subscriber count alone.
2. Skip to content **dormant.** Karpathy, Umar Jamil, Jay Alammar, Aladdin Persson, AI Coffee Break, James Briggs, Arxiv Insights, Connor Shorten — enormous averages

computed on stale evergreen hits. Reference libraries, not feeds. Plan your learning accordingly: back-catalog time is scheduled study, not passive subscription.

3. **Hype is formulaic but algorithmically rewarded.** Berman, Wes Roth, TheAIGrid, and Shapiro post 🔥 daily and pull 13K–98K averages while running exactly the "SHOCKING / INSANE / is HERE (WOAH)" titles their own communities criticize. The counter-proof that depth *can* scale: Dwarkesh (~309K), bycloud (~83K), MLST (~37K). The audience for substance exists; most creators just don't compete for it.

4. **Agentic coding is the hottest zone on the platform by cadence.** Cole Medin posting 6 hours before my scrape, Nate Herk 11h, Riley Brown 7h, AI Engineer within hours of events — nearly all of Tier 10/11 is 🔥, and everyone is converging on Claude Code / Codex / harness engineering. That's where the frontier is moving, measured by creator behavior rather than anyone's opinion.

5. **ComfyUI/SD is quietly cooling at the top.** Olivio has slowed to 🐢, Scott Detweiler has gone 💤 — while Nerdy Rodent, Sebastian Kamph, and the now-verified Latent Vision remain the active reliable core. Content ecosystems have lifecycles; this one has plateaued.

6. **Robotics skews entertainment-heavy.** Reeves (8M avg), Stuff Made Here (5M), Code Bullet (2.4M) dominate reach — but for rigorous robotics you want Skyentific, Articulated Robotics, and Paul McWhorter's "AI on the Edge" at 1/1000th the views.

7. **Benchmark numbers are configuration-specific.** The 40-vs-130 tok/s Radeon discrepancy is the emblem: runtime choice (Ollama vs llama.cpp vs vLLM) can 3× a result on identical hardware. Never quote a creator's benchmark as universal — including the ones in this guide.

Final Thought: Curate Like an Engineer

The AI content ecosystem in 2026 has a shape, and once you see it you can't unsee it:

- Skip to content are dormant but immortal — study them like books.
- The **frontier** is covered brilliantly by a handful of channels that prove depth scales.

- The **news** layer is a routing system — let it point you at primary sources, never let it be the source.
- The **hardware** channels are honest but configuration-bound — trust the method, verify the number.
- The **builders** are moving faster than any other zone — that velocity *is* the signal about where this field is going.

You wouldn't build a production system on unverified claims and vanity metrics. Don't build your information diet on them either.

Start with three:

- **3Blue1Brown** — so you understand what's underneath
- **AI Explained** — so you know what actually happened this week
- **One channel from your work-specific stack** in Part 8 — so it compounds

Then prune quarterly, the way this directory does: check the recency, check the real engagement, and unsubscribe without sentiment.

Methodology

Scope, Dates, and Data Collection

- **Evidence snapshot:** July 12–13, 2026.
- **Core scrape:** 68 channels scraped live, then merged with three research-addendum passes from the same week.
- **Directory size:** roughly 185 distinct channels across 24 active thematic tiers, plus the separately labeled Tier 18 legacy section, a dormant/dead list, a blacklist, and a not-admitted log.
- **Recent-view window:** approximately the latest 15 public uploads per channel. This
Skip to content
verage.
- subscriber count, recent view minimum/maximum/average, latest-upload themes, upload recency, source verification, and attributed Reddit reputation

evidence.

- **Reconciliation rule:** conflicting stats resolve to the freshest scrape. The original values remain disclosed in the reconciliation table.
- **Alias rule:** handles and creator aliases were de-duplicated when the supplied corpus established they were the same channel.
- **Editorial treatment:** the user-supplied corpus is authoritative for its stated scrape dates. This publication pass reorganized, copy-edited, and localized the corpus; it did not run a new web scrape.

Inclusion and Exclusion Rules

A channel was included when it had a recognizable AI, ML, data-science, local-inference, agentic-engineering, creative-AI, robotics, semiconductor, vendor, or adjacent practitioner role and enough current or evergreen evidence to justify a place in the directory. Some channels appear in more than one tier when their work genuinely spans categories; they still count as one distinct channel.

Candidates found in Reddit recommendations were not admitted when the corpus could not recover a canonical YouTube page, establish a stable AI-first identity, or verify enough current activity to classify them honestly. The not-admitted table is part of the evidence trail, not an invitation to infer that those creators are low quality.

Coverage and Honest Gaps

- **Still unverified after repeated provider failures:** Level1Techs, Venelin Valkov, Donato Capitella's video list (subscriber count confirmed at 99.7K), BMad Code, SECourses, Benji's AI Playground, Peter Yang, Articulated Robotics, This Week in AI, Practical AI, Michael Bronstein, Jeff Heaton, Henry AI Lab, plus metrics for Neural Breakdown with AVB, Gal Lahat, Reducible, Brandon Foltz, and zedstatistics.
- **Partial rendering:** StatQuest's average is based on only three recent view counts because the rest did not render during the scrape.
- **ce:** the supplied corpus includes attributed quotations, usernames, and upvote counts, but not exact thread URLs for every quotation. This report preserves those attributions and does not fabricate missing links.

Skip to content

- **Freshness decays:** every 🔥, ✅, 🐢, and 💤 marker describes July 12–13, 2026. A channel can move from active to dormant within a quarter; the ranking method should outlive any individual grade.
- **Benchmark comparability:** local-model performance depends on runtime, quantization, context length, drivers, kernels, and other configuration details. The reported values describe creators' tested setups, not universal hardware guarantees.
- **Signal grades are editorial:** 🟢, 🟡, and 🔴 synthesize the supplied research and Reddit reputation record. They are not objective safety or quality certifications.

Reproducibility Notes

The corpus reports ten stat conflicts reconciled to the freshest scrape, preserves every unverifiable metric with ❓, and keeps the graveyard separate from the blacklist: dormancy, low current utility, hype, and active trust concerns are different findings. The source also distinguishes 24 active thematic tiers from the Tier 18 legacy/sparse section placed inside the graveyard; this report preserves that original organization instead of forcing a misleading renumbering.

The directory should be refreshed quarterly by re-running the recent-upload scrape, checking alias collisions, re-evaluating recency grades, and reviewing whether formerly unverified channels now render. Subscriber counts should remain context, not the ranking target.

Sources

The load-bearing dataset for this report is the **user-provided July 12–13, 2026 live scrape and three same-week research synthesis passes**. The frontmatter source list points to direct channel pages across the main tiers. The following links are the principal public reference points named in the corpus; the complete channel catalog appears in the tier tables above.

Skip to content

Foundations and Implementation

- 3Blue1Brown
- StatQuest
- Welch Labs
- Andrej Karpathy
- Sebastian Raschka
- Umar Jamil
- DeepLearning.AI
- MIT 6.S191 / Alexander Amini

Frontier, News, and Interviews

- bycloud
- Yannic Kilcher
- Machine Learning Street Talk
- Dwarkesh Patel
- AI Explained
- Theo – t3.gg
- AI Daily Brief
- Hard Fork

Hardware, Builders, and Creative Systems

- Alex Ziskind
- Donato Capitella
- Token Chaser
- Cole Medin
- IndyDevDan
- Skip to content
- AI Engineer

- Latent Vision
- Mickmumpitz
- Asianometry

Primary Vendor Channels

- OpenAI
- Google DeepMind
- Anthropic
- Hugging Face

Reddit Reputation Sources Named in the Corpus

- r/learnmachinelearning
- r/LocalLLaMA
- r/MachineLearning
- r/ClaudeAI

These communities form the reputational backbone of the supplied grades and quotations. Exact thread URLs were not included for every quote, so they are not invented here.

Sources & further reading

1. 3Blue1Brown
2. Andrej Karpathy
3. Sebastian Raschka
4. DeepLearning.AI
5. bycloud
6. Skip to content g Street Talk
7. Dwarkesh Patel

8. AI Explained
9. Theo – t3.gg
10. Alex Ziskind
11. Donato Capitella
12. Cole Medin
13. IndyDevDan
14. Sam Witteveen
15. AI Engineer
16. Latent Vision
17. Anthropic
18. Google DeepMind
19. OpenAI
20. Asianometry

Questions & answers

- What is the best AI YouTube channel to follow in 2026?
- How were the channels ranked?
- Which channels are best for AI coding and agents?
- Can I trust local-LLM hardware benchmarks from YouTube?

✳

See you down the next rabbit hole.

More notes →