


PHASE 1 · FOUNDATIONS / 45 MINUTES

Every AI Model Explained

Why there is no single "best" model - and how do you choose the right kind of tool?

**BIG QUESTION**

How are current LLMs different?

BY THE END, I CAN...

Sort models into useful families, compare quality, speed, cost, and control, and justify a model choice for a real task.

MISSION OUTPUT

Model map + model-match recommendation + exit sentence.

LEARN THIS

Large language models (LLMs) are machine-learning models trained on lots of examples of language. They predict useful next tokens, but different models are trained, tuned, hosted, and priced differently. Think in **families**, not a permanent leaderboard: **flagship** models aim for the hardest reasoning and coding; **workhorse** or mid-tier models handle most daily tasks; **light** models trade some capability for speed and low cost. Some are **open-weight** so people can download and adapt them; others are hosted as closed services. **Multimodal** models work with more than text, while specialized models focus on coding, images, video, audio, search, or retrieval. A model's answer is still a draft to test - not proof that it is correct.

YOUR 45-MINUTE PLAN

0-5	Launch: what AI tools have you tried?
5-15	Watch the model map; capture four categories
15-25	Mini-lesson: capability, modality, and trade-offs
25-40	Model-match lab: choose a tool for four jobs
40-45	Explain your choice; complete the exit ticket

SIX STOPS ON THE MODEL MAP

Flagship Top capability for difficult reasoning, long projects, and complex coding. Usually costs more or runs slower.	Workhorse Strong everyday models: a useful balance of quality, speed, cost, and availability.	Light Small and fast. Great for classification, short drafts, simple chat, and high-volume automation.
Open-weight Weights can be downloaded, run, or adapted. Local use can add control, but needs hardware and care.	Multimodal Accepts or creates more than text: images, audio, video, or mixed inputs. Check which modes are supported.	Specialized Coding, search/RAG, enterprise workflows, agents, image, video, and audio models are built for specific jobs.

CHOOSE BY JOB, NOT HYPE

- Name the job:** What must the model receive, produce, and get right? A current handbook answer is different from a poem or a batch of labels.
- Score the trade-offs:** rate the need for quality, speed, cost, modality, context length, privacy, and tool access from low to high. There is no "best" score for every task.
- Test the same prompt:** try two reasonable models with safe, non-private test data. Compare the outputs, sources, latency, and failure modes; keep a human in charge of the final decision.

FOR A CURRENT RESEARCH QUESTION, I WOULD CHOOSE...

MY EVIDENCE FOR THAT CHOICE IS...

MODEL SCORECARD

Quality	How accurate, capable, or nuanced does the task need to be?
Speed	Does a person wait, or must the system handle many requests?
Cost	Is the task occasional, or repeated thousands of times?
Control	Do privacy, hosting, customization, or audit needs matter?

FAMILIES IN THE VIDEO

Examples named or shown include Claude, GPT, Gemini, Grok, Kimi, Qwen, GLM, Llama, Mistral, MiniMax, Momo, Cohere, Amazon Nova, and specialized image, video, and audio tools. Treat these as a time-stamped snapshot. A company's "flagship" label is marketing; compare the actual task, access, and evidence.

WORDS TO KNOW

LLM · model family · flagship · workhorse · light model · open-weight · multimodal · context window · inference · RAG

QUICK CHECK

Match each job to a model category and explain one trade-off: (1) sort 10,000 short messages, (2) solve a hard coding bug, (3) answer from a current school handbook, (4) run offline on a family computer.

SOURCE VIDEO / REFERENCE

<https://www.youtube.com/watch?v=-8pJvYNcX4>
Every AI Model Explained In 20 Minutes (Update)
Examples and rankings change quickly; re-check current model details before choosing a tool.

SAFETY / ETHICS

Do not paste passwords, school records, private photos, medical details, or someone else's data into an AI tool without teacher approval. Open-weight does not automatically mean safe or private. Check licenses, sources, bias, and outputs. Disclose AI assistance when required.