AI STUDIO · CHAPTER 1

Pick your AI

The big AI models are all strong now. What changes is what their rules allow, what they cost and whether you like how they answer. Test before you trust.

A1	5 min	4
ENTRY	READING	SECTIONS



EVIDENCE ◆ Verified ◆ Vivian's field notes ◆ Unverified

IN 30 SECONDS

01 The big models are close in power. Pick the one whose answers you'd actually post, after testing it.

02 Read a provider's rules before you use it for adult work. Most of them forbid explicit sexual text.

03 Prices go from a few cents to tens of dollars per million tokens, so the same job can cost very different amounts.

04 Two versions of the same model can behave differently. Test the one you'll really use.

05 Never paste your legal name, your payment details or a sub's private messages into a chat.

YOU WILL LEARN

- ✓ Why testing beats any ranking.
- ✓ What each big provider's rules say about adult text.
- ✓ What the main models cost, by the million tokens and by the month.
- ✓ What to keep out of an AI chat.

PART 01

Test before you trust

Every big model can write, plan and code well enough for 2DFD work. What changes is how you use it and whether you like its answers. ◆ A ranking

won’t tell you that, so test.

A test you can run in one evening

1. Write three requests you really need: a short caption, a calm reply to a sub who asks for more, and a small piece of code.

2. Send the same three to two or three models, word for word.

3. Read the answers the next day, side by side, without looking at which model wrote which.

4. Keep the model whose answers you’d post as they are. Note what it refused, too.

TRY IT

Same request, two answers

Tap every message that shows a red flag. There are two.

Example conversation. Invented for this exercise.

YOU

Write a one-line caption for a drawn persona posting a new wishlist item. Teasing, not explicit.

THEM

New on my list. You know where it is.

Red flag 1. Short, in character, and it says nothing a platform would flag.

YOU

Same request, second model.

THEM

Exciting news, everyone. Check out my amazing new wishlist item and show some love.

Red flag 2. Same job, but it sounds like an ad. You'd rewrite it, so this model costs you time.

The test also shows you the refusals. A model that refuses half of your normal work isn't the right tool, however good its scores are. ◆

What the rules say

Each provider publishes a usage policy, and it applies to everything you send. Read it before you use a model for adult work, because breaking it can close your account. ◆

A DOCUMENTED ADULT MODE All No Not for text Not found

MODEL ↕	WHAT ITS PROVIDER'S RULES SAY ABOUT SEXUAL TEXT	A DOCUMENTED ADULT MODE ↕
GPT-6 Luna (OpenAI) ☐ compare	By default, no erotica. OpenAI says it's exploring an adult mode	No
Claude Sonnet 5.5, Opus 4.6 and 4.8 (Anthropic) ☐ compare	No sexually explicit content, no fetish content, no erotic chats	No
Grok 4.7 (xAI) ☐ compare	No pornographic use of real people's likeness. Its terms mention 18+ adult content	Not for text
Kimi K3.1 Pro (Moonshot) ☐ compare	Its consumer terms forbid obscenity and pornography	No
GLM-5.3 (Z.ai) ☐ compare	Forbids sexually explicit content and even suggestive content	No
Qwen3.8-Max (Alibaba) ☐ compare	The rules read ban content that harms minors and don't address adult text	Not found

	WHAT ITS PROVIDER'S RULES SAY ABOUT SEXUAL TEXT	
DeepSeek ☐ compare	No pornographic content and no sexual chatbots	No

OpenAI's default rules say its model should not generate erotica. ◆ OpenAI also writes that it's exploring a way to allow it in age-appropriate contexts. ◆ Anthropic's policy forbids content about sexual fetishes or fantasies and erotic chats, for every Claude model. ◆ xAI's policy bans depicting real people's likenesses in a pornographic manner. ◆ Its consumer terms mention 18+ adult content and a setting to limit it. ◆ Z.ai's terms forbid content that is sexually explicit or suggestive. ◆ DeepSeek's terms forbid sexually explicit content and sexual chatbots. ◆ None of these rules names findom. ◆

VIVIAN'S FIELD NOTES

In my experience, Grok is the most open of these to adult and suggestive work, and its terms are the only ones here that mention 18+ content. I also find that two versions of the same model don't apply the same rules with the same strictness. For apps that are more about control, like a hub that sends example requests, Opus 4.6 and 4.8 refused less than the newest versions. That's my own test, not something any provider says. Run yours.

◆

Don’t try to talk a model out of its rules. Accounts that do get banned, and you lose your history with them. ◆ If a model refuses work you need, switch to one whose rules allow it.

PART 03

What it costs

There are two ways to pay. A monthly plan gives you the chat app. The API charges by the token, which is roughly a piece of a word, and it’s what you use when a tool you build talks to the model. ◆

MODEL ↕	API INPUT, PER MILLION TOKENS ↕	API OUTPUT, PER MILLION TOKENS ↕
GPT-6 Luna ☐ compare	USD 0.10	USD 0.50
Claude Sonnet 5.5 ☐ compare	USD 2	USD 10
Claude Opus 4.6 or 4.8 ☐ compare	USD 5	USD 25
Grok 4.7 ☐ compare	USD 2.00	USD 6.00
GLM-5.3 ☐ compare	USD 1.4	USD 4.4
Qwen3.8-Max ☐ compare	USD 2	USD 6
DeepSeek V4 Pro, peak hours ☐ compare	USD 1.32	USD 3.96

GPT-6 Luna costs USD 0.10 per million input tokens and USD 0.50 per million output tokens on OpenAI's API. ◆ Claude Sonnet 5.5 costs USD 2 in and USD

10 out, and Opus 4.6 and 4.8 cost USD 5 in and USD 25 out. ◆ Grok 4.7 costs USD 2.00 in and USD 6.00 out for prompts under 200,000 tokens. ◆ GLM-5.3 costs USD 1.4 in and USD 4.4 out. ◆ Qwen3.8-Max costs USD 2 in and USD 6 out in Alibaba's international region. ◆ DeepSeek charges half price off-peak. ◆ Prices were read on 2026-11-07 and change often. ◆

A caption is a few dozen tokens, so a month of captions costs cents on the API. Code is where tokens add up, because the model reads your files again on every request. ◆ The budget planner in this section turns this into a monthly figure for your own use.

PART 04

What to keep out of the chat

What you paste into a chat leaves your computer. ◆ Keep these out:

Never paste

1. Your legal name, your address, or a document that shows them.
2. Payment details, passwords, keys or login codes.
3. A sub's private messages, real name or money details.
4. Screenshots that show your other accounts or your desktop.

Use the persona's name and made-up examples instead. The model doesn't need the real thing to help. ◆ See Stay anonymous.

TRY IT

Your first model

You want an AI to help with captions and a small wishlist tool. A friend says one model is "the best".

A Pay for a year of the model your friend named

Not this one. A ranking doesn't know your work. You could pay a year for answers you don't like.

B Run the same three requests through two or three models and keep the one you'd post

The right answer. Right. Testing costs an evening and shows you the refusals too.

C Use whichever model refuses nothing

Not this one. A model that ignores its own rules today can close your account tomorrow, with your history.

REMEMBER IT

- The big models are close. Test, then pick the answers you like.
- Read the provider's rules before adult work.
- Prices differ a lot: check the API price for anything you build.
- Never paste real names, payment details or a sub's private data.

SOURCES AND EVIDENCE

Check it

Every claim with its evidence level and its sources.

01 Every big model can write, plan and code well enough for 2DFD work.

02 What changes is how you use it and whether you like its answers.

03 A model that refuses half of your normal work isn't the right tool, however good its scores are.

04 Read it before you use a model for adult work, because breaking it can close your account.

05 OpenAI's default rules say its model should not generate erotica.

- **OpenAI Model Spec (2026-10-29)** (https://model-spec.openai.com/2026-10-29.html)

Accessed 7 November 2026

06 OpenAI also writes that it's exploring a way to allow it in age-appropriate contexts.

- **OpenAI Model Spec (2026-10-29)** (https://model-spec.openai.com/2026-10-29.html)

Accessed 7 November 2026

07 Anthropic's policy forbids content about sexual fetishes or fantasies and erotic chats, for every Claude model.

- **Anthropic Usage Policy** (https://www.anthropic.com/legal/aup)

Accessed 7 November 2026

08 XAI's policy bans depicting real people's likenesses in a pornographic manner.

- **xAI Acceptable Use Policy** (https://x.ai/legal/acceptable-use-policy)

Accessed 7 November 2026

09 Its consumer terms mention 18+ adult content and a setting to limit it.

- **xAI Terms of Service (consumer)** (https://x.ai/legal/terms-of-service)

Accessed 7 November 2026

10 Z.ai's terms forbid content that is sexually explicit or suggestive.

- **Z.ai Terms of Use** (https://docs.z.ai/legal-agreement/terms-of-use)

Accessed 7 November 2026

11 DeepSeek's terms forbid sexually explicit content and sexual chatbots.

◆

DeepSeek Terms of Use (https://cdn.deepseek.com/policies/en-US/deepseek-terms-of-use.html)

Accessed 7 November 2026

12 None of these rules names findom.

◆

13 In my experience, Grok is the most open of these to adult and suggestive work, and its terms are the only ones here that mention 18+ content. I also find that two versions of the same model don't apply the same rules with the same strictness. For apps that are more about control, like a hub that sends example requests, Opus 4.6 and 4.8 refused less than the newest versions. That's my own test, not something any provider says. Run yours.

◆

14 Accounts that do get banned, and you lose your history with them.

◆

15 The API charges by the token, which is roughly a piece of a word, and it's what you use when a tool you build talks to the model.

◆

16 GPT-6 Luna costs USD 0.10 per million input tokens and USD 0.50 per million output tokens on OpenAI's API.

◆

GPT-6 Luna model page (https://developers.openai.com/api/docs/models/gpt-6-luna)

Accessed 7 November 2026

17 Claude Sonnet 5.5 costs USD 2 in and USD 10 out, and Opus 4.6 and 4.8 cost USD 5 in and USD 25 out.

◆

Claude API pricing (https://platform.claude.com/docs/en/about-claude/pricing)

Accessed 7 November 2026

18 Grok 4.7 costs USD 2.00 in and USD 6.00 out for prompts under 200,000 tokens.

- **xAI models and pricing** (https://docs.x.ai/developers/models)

Accessed 7 November 2026

19 GLM-5.3 costs USD 1.4 in and USD 4.4 out.

- **Z.ai pricing** (https://docs.z.ai/guides/overview/pricing)

Accessed 7 November 2026

20 Qwen3.8-Max costs USD 2 in and USD 6 out in Alibaba's international region.

- **Alibaba Cloud Model Studio pricing** (https://www.alibabacloud.com/help/en/model-studio/model-pricing)

Accessed 7 November 2026

21 DeepSeek charges half price off-peak.

- **DeepSeek API pricing** (https://api-docs.deepseek.com/quick_start/pricing/)

Accessed 7 November 2026

22 Prices were read on 2026-11-07 and change often.

23 A caption is a few dozen tokens, so a month of captions costs cents on the API.

24 Code is where tokens add up, because the model reads your files again on every request.

25 What you paste into a chat leaves your computer.

◆

26 The model doesn't need the real thing to help.

◆