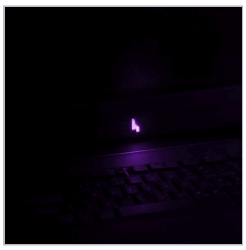
AI STUDIO · CHAPTER 2

~$ open captions-with-ai

Captions with AI

A model can draft captions, post text and replies in your persona's voice once you've shown it that voice. Then you edit, so a person can hear a person in it.

ENTRY A2 CATEGORY AI STUDIO CHECKED 7 NOVEMBER 2026 READING 6 MIN
SECTIONS 7



ls cover/ A2.webp

EVIDENCE ◆ Verified ◆ Vivian's field notes ◆ Unverified

IN 30 SECONDS

01 Give the model a short voice sheet and three real captions before you ask for anything.

02 Ask for several drafts, keep one, and rewrite it in your own words.

03 Cut what gives a machine away: ad words, stray hashtags, every line built the same way.

04 Suggestive captions are fine on most platforms. Explicit content and selling for sends are where the rules bite.

05 A sub's name, messages and amounts never go into a chat with a model.

YOU WILL LEARN

- ✓ What goes on a voice sheet and how to use it.
- ✓ What a weak and a good request to a model look like.
- ✓ How to edit a draft so it sounds human.
- ✓ What the providers' and the platforms' rules mean for teasing captions.
- ✓ How to get help with a reply without handing over a sub's private data.

> *A model can write a hundred captions a minute. None of them is yours until you've cut it down and said it your way.*
>
> — VIVIAN

~/01-show-it-the-voice-first

Show it the voice first

With no instructions, a model writes like an ad. So before you ask for a caption, give it a voice sheet: a few lines on who the persona is and how she talks. In my experience, ten lines work better than two pages, because the model follows a short sheet closely and a long one loosely. ◆

```
# A voice sheet in five parts
```

a-voice-sheet-in-five-parts.md · 5 lines copy

```
1 Who she is, in one line: her name, her look, the kind of dom
  she plays.
2 How she talks: how long her sentences are, words she uses, words
  she never uses.
3 How she treats the reader: teasing, cold, warm or mocking.
4 Three captions you've posted and liked, pasted as they are.
5 What she never writes: explicit sex, real names, promises
  about money.
```

The examples do more work than the description. A model picks up rhythm from three real captions far better than from words like “playful” or “dominant”. Swap the examples every few weeks, or every draft starts to sound like the same three posts. ◆ Most of the sheet can come straight from your persona file, as Build a persona explains.

~/02-ask-for-ten-keep-one

Ask for ten, keep one

Ask for ten drafts at once, each under a set length, and pick the best. Choosing is faster than fixing one bad draft. Say what the post is for: a new picture, a wishlist item, a tool launch, a quiet day. ◆

TRY IT

A weak request

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

Example conversation. Invented for this exercise.

YOU

Write a sexy findom caption for my new pic.

THEM

Ready to spoil your Goddess? My new pic is here and it's waiting for your tribute. Don't keep me waiting.

Red flag 1. Generic and pushy, and it could belong to any account. "Sexy findom" told the model nothing about who's talking.

YOU

Make it more dominant.

THEM

Kneel. Pay. Obey. You know your place.

Red flag 2. More orders, same emptiness. Adjectives don't give a model a voice.

TRY IT

A good request

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

Example conversation. Invented for this exercise.

YOU

Here's my voice sheet and three captions I've posted. Write ten captions for a new picture: she's on a throne, looking down at the viewer. Under twelve words each, teasing, never explicit. Leave out hashtags, emoji and anything about money.

THEM

Comfortable down there? Good. Stay.

Red flag 1. Short and in voice. Keeping money out stops the caption turning into a sales line.

I didn't pick this chair for you. You can still look.

Red flag 2. Close. One word changed and it's yours.

YOU

Number 2 is close. Make it colder and shorter.

THEM

Not your chair. Look anyway.

Red flag 3. The second round is where it gets good. Narrow it down instead of starting again.

Make it sound like you

Every draft gets a human pass before it goes out. Read it out loud, and if you wouldn't say it, change it. ◆ Then check it still says one clear thing, as Posting and workflows explains.

What gives a machine away

RED FLAG 1
Ad words: "exciting", "amazing", "don't miss", "check out".

RED FLAG 2
Hashtags and emoji you didn't ask for.

RED FLAG 3
Every line built the same way, often in threes.

RED FLAG 4
A question tacked onto the end of every post.

RED FLAG 5
A long word where she'd use a short one.

RED FLAG 6
The same opener day after day.

Keep a note of the changes you make most often. After a month, add them to the voice sheet, so the model stops making the same mistakes. ◆

OpenAI's terms list presenting output as human-generated when it wasn't among the things users may not do. ◆ Read your own provider's terms on this before you post drafts as they come.

~/04-what-the-providers-allow

What the providers allow

A teasing caption meets the providers' rules on one side and the platforms' rules on the other. ◆ Pick your AI has the full table of providers. For captions, a few lines matter most.

Z.ai's terms forbid suggestive content as well as explicit content, so even a teasing caption is outside them. ◆ Anthropic forbids content related to sexual fetishes or fantasies, for every Claude model. ◆ Kimi's consumer terms forbid content that promotes obscenity or pornography. ◆ xAI's consumer FAQ says Grok may answer with sexual situations, among other things. ◆ On xAI's API, a developer can tell the model to decline sexual innuendo, and leaving that setting out keeps the default behaviour. ◆ No policy read names findom, so whether a money-teasing caption counts as fetish content is the provider's call. ◆

If a model won't write the captions you need, use one whose rules allow them, or write that kind yourself. Prompting around a model's rules can cost you the account. ◆

~/05-what-platforms-flag

What platforms flag

Suggestive and explicit are two different lines, and most platforms draw theirs between them. ◆ Throne allows suggestive but not explicit content on a wishlist, and applies its rules to drawn art as it does to photos. ◆ X allows adult content that's labelled and not prominently displayed, including anime and AI-generated art. ◆

The warnings X describes go on photos and videos, either account-wide or one post at a time. ◆ What X's classifier reads in a caption's text isn't published. ◆ The full method for marking posts is in Posting and workflows.

The words can still break a rule on their own. X's adult content rules bar promoting or soliciting sexual services. ◆ Throne bars profile text that offers content or services in exchange for gifts. ◆ A caption that sells a picture for a send can fall under those lines, so read them before you write one. ◆

~/06-keep-the-sub-out-of-it

Keep the sub out of it

Replies are where AI help gets risky, because the material belongs to someone else. A sub's messages, name, handle, photos and amounts are his. Pasting them into a chat hands them to a company he never agreed to. ◆

```
# Help with a reply, without his data
```

help-with-a-reply-without-his.md · 5 lines copy

1. Describe the situation in your own words: “a sub asks for a bigger task than the one agreed”.
2. Write “he” or “the sub” instead of his name or handle.
3. Leave out amounts, or round them.
4. Never paste screenshots of the chat.
5. Ask for three short replies in the persona’s voice, then rewrite the one you pick.

Write prices, totals and dates yourself. A model can get a number wrong, and in a money message the mistake is yours. ◆

Some replies shouldn’t be drafted by a model at all. If a sub says money is short, mentions debt or sounds unwell, stop the money part and answer him yourself, in plain words. ◆ See When a sub is in trouble.

~/07-what-it-costs

What it costs

Captions don’t need the API. A chat plan is enough, and the providers below also have a free tier. ◆

◆	Price read on 2026-11-07
ChatGPT Go and Plus	Shown in your local currency
Claude Pro	USD 20 a month, billed monthly
Grok Plus	Shown in your local currency

Claude Pro costs USD 20 when billed monthly. ◆ ChatGPT and Grok show their plans in the visitor's own currency, so read the price on their pricing pages. ◆ Prices depend on your country and change often. ◆ If you build a tool that writes captions for you, the API prices are in Pick your AI.

TRY IT

A long message

A sub sends a long message about a task he finished, with his real first name in it, and asks what comes next. You want help with the answer.

A Paste his message into the chat and ask for a reply in your voice

Not this one. His words and his name go to a provider he never agreed to.

B Sum it up in two lines without his name, and ask for three drafts

The right answer. Right. The model gets what it needs, and he stays private.

C Let a tool answer him directly with whatever the model writes

Not this one. Nobody reads it before he does. A wrong number or a wrong tone goes out under your name.

REMEMBER IT

- Give the model a short voice sheet and three real captions first.
- Ask for ten drafts, keep one, then rewrite it.
- Cut ad words, stray hashtags and lines that all sound the same.
- Teasing is fine on most platforms. Explicit content and selling for sends are where the rules bite.
- A sub's messages, name and amounts stay out of the chat.

SOURCES AND EVIDENCE

Check it

Every claim with its evidence level and its sources.

01 With no instructions, a model writes like an ad.

◆

02 So before you ask for a caption, give it a voice sheet: a few lines on who the persona is and how she talks.

◆

03 In my experience, ten lines work better than two pages, because the model follows a short sheet closely and a long one loosely.

◆

04 The examples do more work than the description.

◆

05 A model picks up rhythm from three real captions far better than from words like “playful” or “dominant”.

◆

06 Swap the examples every few weeks, or every draft starts to sound like the same three posts.

◆

07 Ask for ten drafts at once, each under a set length, and pick the best.

◆

08 Choosing is faster than fixing one bad draft.

◆

09 Say what the post is for: a new picture, a wishlist item, a tool launch, a quiet day.

10 Every draft gets a human pass before it goes out.

11 Read it out loud, and if you wouldn't say it, change it.

12 After a month, add them to the voice sheet, so the model stops making the same mistakes.

13 OpenAI's terms list presenting output as human-generated when it wasn't among the things users may not do.

- **OpenAI Terms of Use (outside the US and EEA)** (https://openai.com/policies/row-terms-of-use/)

Accessed 7 November 2026

14 A teasing caption meets the providers' rules on one side and the platforms' rules on the other.

15 Z.ai's terms forbid suggestive content as well as explicit content, so even a teasing caption is outside them.

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

Accessed 7 November 2026

16 Anthropic forbids content related to sexual fetishes or fantasies, for every Claude model.

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

Accessed 7 November 2026

17 Kimi's consumer terms forbid content that promotes obscenity or pornography.

◆ **Kimi Terms of Service (international, consumer)** (https://www.kimi.ai/user/agreement/modelUse)

Accessed 7 November 2026

18 XAI's consumer FAQ says Grok may answer with sexual situations, among other things.

◆ **xAI consumer FAQ** (https://x.ai/legal/faq) Accessed 7 November 2026

19 On xAI's API, a developer can tell the model to decline sexual innuendo, and leaving that setting out keeps the default behaviour.

◆ **xAI API docs: Custom moderation controls** (https://docs.x.ai/developers/advanced-api-usage/custom-moderation-controls)

Accessed 7 November 2026

20 No policy read names findom, so whether a money-teasing caption counts as fetish content is the provider's call.

◆

21 If a model won't write the captions you need, use one whose rules allow them, or write that kind yourself.

◆

22 Prompting around a model's rules can cost you the account.

◆

23 Suggestive and explicit are two different lines, and most platforms draw theirs between them.

◆

24 Throne allows suggestive but not explicit content on a wishlist, and applies its rules to drawn art as it does to photos.

◆

What's allowed in my profile text and images? (https://help.throne.com/en/articles/15810290-what-s-allowed-in-my-profile-text-and-images)

Archive (https://web.archive.org/web/20261026034619/https://help.throne.com/en/articles/15810290-what-s-allowed-in-my-profile-text-and-images)

Accessed 3 November 2026

25 X allows adult content that's labelled and not prominently displayed, including anime and AI-generated art.

◆

Adult Content Policy (https://help.x.com/en/rules-and-policies/adult-content)

Archive (https://web.archive.org/web/20261002155114/https://help.x.com/en/rules-and-policies/adult-content)

Accessed 3 November 2026

26 The warnings X describes go on photos and videos, either account-wide or one post at a time.

◆

Your media settings (https://help.x.com/en/rules-and-policies/media-settings)

Archive (https://web.archive.org/web/20261016111110/https://help.x.com/en/rules-and-policies/media-settings)

Accessed 3 November 2026

27 What X's classifier reads in a caption's text isn't published.

◆

28 X's adult content rules bar promoting or soliciting sexual services.

◆

Adult Content Policy (https://help.x.com/en/rules-and-policies/adult-content)

Archive (https://web.archive.org/web/20261002155114/https://help.x.com/en/rules-and-policies/adult-content)

Accessed 3 November 2026

29 Throne bars profile text that offers content or services in exchange for gifts.

◆

What's allowed in my profile text and images? (https://help.throne.com/en/articles/15810290-what-s-allowed-in-my-profile-text-and-images)

Archive (https://web.archive.org/web/20261026034619/https://help.throne.com/en/articles/15810290-what-s-allowed-in-my-profile-text-and-images)

Accessed 3 November 2026

30 A caption that sells a picture for a send can fall under those lines, so read them before you write one.

◆

31 Replies are where AI help gets risky, because the material belongs to someone else.

◆

32 Pasting them into a chat hands them to a company he never agreed to.

◆

33 A model can get a number wrong, and in a money message the mistake is yours.

◆

34 If a sub says money is short, mentions debt or sounds unwell, stop the money part and answer him yourself, in plain words.

◆

35 A chat plan is enough, and the providers below also have a free tier.

◆

36 Claude Pro costs USD 20 when billed monthly.

◆ **Claude plans and pricing** (https://claude.com/pricing)

Accessed 7 November 2026

37 ChatGPT and Grok show their plans in the visitor's own currency, so read the price on their pricing pages.

◆

38 Prices depend on your country and change often.

◆