

CaptionLab

Windows Install Guide — get auto subtitles working in Premiere Pro

CaptionLab installs with a small installer script — no admin rights needed, nothing to configure by hand.

What you need

- Windows 10 or 11 (64-bit)
- Adobe Premiere Pro 2022 or later
- ~3.5 GB free disk space for first-time setup (CPU-only) — **~5.5 GB** if you have an NVIDIA GPU (adds CUDA runtime libraries for faster transcription)
- More disk space only if you choose to download additional AI models later from the **Models** tab (see below) — each one is optional, from 480 MB to 3.1 GB
- Internet connection for one-time setup and model downloads only (transcription itself is fully offline)

Install in 4 steps

⚠ Installed CaptionLab before? (upgraders only)

If you've ever installed any version of CaptionLab on this PC — especially via a different tool (a ZXP installer, or a manual copy) — first double-click `Uninstall CaptionLab.bat` (included in the zip) and let it finish. It removes every old copy so they don't conflict with the new one. **Brand-new users can skip this** — go straight to step 1.

1 Unzip the file you downloaded — extract everything to a folder first. Don't run anything from inside the zip viewer.

2 Quit Premiere Pro completely if it's running.

3 Double-click **Install CaptionLab.bat** in the extracted folder. It installs the extension, enables Premiere's extension support, and cleans up any old install automatically — no Terminal, no typing.

4 Open Premiere Pro. Go to `Window → Extensions → CaptionLab`, then click **"Install AI Engine"** on first launch.

CaptionLab automatically downloads Python, ffmpeg, and the default Whisper AI model — entirely inside the extension. Takes about 15–30 minutes depending on your internet connection. If you have an NVIDIA GPU, it also installs the CUDA runtime libraries needed to use it.

GPU vs CPU speed

Transcription runs faster with an NVIDIA GPU. Below is an approximate real-time factor (how many minutes of audio get transcribed per minute of processing) — **estimated**, based on one measured card, not benchmarked on every model listed. Faster models (see the Models tab section below) speed this up further on any hardware.

Hardware	Approx. speed
No GPU (CPU only)	~0.5–1x realtime
GTX 1660 Super (no tensor cores)	~2–2.5x
RTX 2060 / 3050	~3.5–4x
RTX 2070 Super (<i>measured</i>)	~4.7–4.9x
RTX 3060	~4–4.5x
RTX 3060 Ti / 3070	~6–7x
RTX 3080	~9–10x
RTX 3090 / 4070	~10–11x
RTX 4080	~13–14x
RTX 4090	~18–20x

Laptop GPUs typically run at ~70–80% of the desktop figure above due to power/thermal limits. Without a compatible GPU, or if the CUDA libraries aren't available, CaptionLab automatically falls back to CPU — it always works, just slower.

Start your free trial

You don't have to buy CaptionLab to use it. A **free 3-day trial** unlocks the complete plugin — every model, every language, real subtitles in your timeline. No card, no account.

- 1 Open **CaptionLab** in Premiere Pro and enter your **email address**.
- 2 Check your inbox — a **trial key** arrives within a minute (look in spam if it doesn't).
- 3 Paste the key into the panel and click **Start my free trial**.
- 4 Click **Install AI Engine**. This downloads the transcription engine and a ~3 GB AI model — by far the longest step, so start it and go do something else.
- 5 When it finishes, select a clip and click **Generate Subtitles**. **This is the moment your 3 days start** — everything before it is free.

Your 3 days start at your first transcription

Not at download, and not when you enter the key. Installing the engine and downloading an AI model can take an evening — none of that comes out of your trial. Take your time.

Test your own language first

Accuracy varies a lot between languages. If you work in one Adobe doesn't support, run a short clip of your own footage through it before deciding — that is exactly what the trial is for.

Activate your license

After purchasing CaptionLab, you'll receive an email with your **Order ID** (starts with `txn_`).

- 1 Open **CaptionLab** in Premiere Pro — you'll see the activation screen.
- 2 Enter the **email address** you used to purchase.

3 Enter your **Order ID** (`txn_01...` from your receipt email).

4 Click **Activate**. That's it — CaptionLab is unlocked forever.

 Works offline

After activation, CaptionLab works offline for 14 days before needing a quick online re-check. No subscription, no recurring fees.

What's new in v1.0.5

If you used an earlier version, these are the changes you'll notice most:

- **A free 3-day trial.** The complete plugin, before you pay anything — and the three days start at your first transcription, not at download.
- **A model built for Burmese.** Burmese is one of the languages the standard models handle badly: on real footage they produced the wrong script entirely, or repeated the same syllable for minutes. A Burmese-specific model is now in the **Models** tab and produces real, readable Burmese on the same clips.
- **Installs that repair themselves.** A half-finished download used to be trusted forever — every retry saw the file was "there" and skipped it, so the same broken install came back however many times you clicked. Each piece is now checked to see it actually runs, and re-fetched if not. Stalled downloads time out and retry instead of hanging.

What's new in v1.0.4


If you used an earlier version, these are the changes you'll notice most:

- **Caption timing is measured against the audio.** Caption starts are now checked against the sound itself rather than taken from the model, so a caption no longer appears while the speaker is still silent. This matters most after long pauses, over B-roll with room tone, and after music.
- **Quiet speech is no longer skipped.** Audio recorded well below the level of the rest of a video — a location interview cut against a loud studio voice, for example — used to be dropped silently. It is now levelled before transcription.
- **Captions follow sentence structure.** A caption is never split across a sentence boundary, breaks at commas and clause boundaries, and short binding words are never left stranded on

their own.

- "Fixed words" is now "Max words" — see below.

How to use — Transcribe tab




The Transcribe tab — this is the default screen every time you open the panel

1. Open **Window → Extensions → CaptionLab**
2. **Source** — select a clip on your timeline and choose **Selected Clip**, or pick **Full Sequence** to caption the whole timeline at once (multicam sequences and Merged/Synchronized clips are supported)
3. **Language** — leave on **Auto-detect**, or pick the spoken language from the dropdown (99 languages supported)

4. **Output** — choose **Transcribe** to keep the original language, or → **English** to translate to English
5. **Caption length** — see the "Max words" section below for the short-form video option
6. Click **Generate Subtitles**
7. Drag the resulting `.srt` file from the Project Panel to your timeline

New in v1.0.4 — Caption length: "Max words"



Caption length set to "Max words" — the Words slider appears, here set to 3 words per caption

By default, CaptionLab groups words into captions the same way a human editor would — by sentence structure, punctuation and pauses. That's right for regular horizontal video, but for

vertical short-form video (Shorts, Reels, TikTok-style edits) you usually want short bursts of 1–3 words popping on screen in sync with speech instead.

To use it:

1. Under **Caption length**, click **Max words** instead of Default
2. A **Words** slider appears — drag it to choose the maximum number of words per caption, from **1 to 8**
3. Generate subtitles as usual

It's a ceiling, not a quota. This setting was called "Fixed words" before v1.0.4, which was misleading. If you set 4, a caption will still be 3, 2 or even 1 word whenever that's where the line logically ends — a caption is never split across a sentence boundary, and it breaks at commas and clause boundaries too. Short binding words ("not", "the", "in" and their equivalents in other languages) are never left stranded alone on screen.

1–3 words works best for the fast-paced, word-by-word caption style common on Shorts/Reels/TikTok. Higher values behave closer to a short readable phrase per caption.

Models tab

CaptionLab

EN

Ready

Transcribe

Models

STORAGE

Using 4.1 GB of models on disk

Large-v2

3.1 GB

Current default. Stable, reliable across all 99 languages.

ACTIVE

Delete

Large-v3

3.1 GB

Better on noisy audio & accents. Same speed as v2.

NOT DOWNLOADED

Download

Large-v3-turbo

1.6 GB

~6x faster. Slight quality trade-off on some languages. Not recommended for translation.

DOWNLOADED

Use

Delete

Small

480 MB

Fast, lightweight. Best for weaker machines or quick drafts.

DOWNLOADED

Use

Delete

CaptionLab v1.0.3

Refresh

The Models tab — Large-v2 is active here, Large-v3-turbo and Small are downloaded but not in use, Large-v3 is available to download

CaptionLab now lets you choose which Whisper AI model does the transcribing, right from the panel — no reinstalling, no Terminal. Click the **Models** tab at the top of the panel to open it.

Model	Size	Best for
Large-v2 (default)	3.1 GB	Stable, reliable across all 99 languages — the safest general-purpose choice

Large-v3	3.1 GB	Better on noisy audio and accents, same speed as v2 — try this if v2 struggles with a particular recording
Large-v3-turbo	1.6 GB	~6x faster than the others — great when you need speed and are transcribing (not translating). Slight quality trade-off on some languages, and not recommended for the "→ English" translate mode
Small	480 MB	Fastest and lightest — best for weaker/older machines or quick rough drafts
Myanmar (Burmese)	2.9 GB	Fine-tuned specifically for Burmese. Use this instead of Large-v3 if you work in Burmese — the general models struggle badly with it

What each button does:

- **Download** — downloads that model to your PC. You'll see a live progress bar. This only needs to happen once per model.
- **Use** — switches CaptionLab to transcribe with that model from now on. Only appears once a model is downloaded.
- **Delete** — removes a downloaded model from disk to free up space. You can always download it again later. (The currently active model can't be deleted — switch to another one first.)

The **ACTIVE** badge shows which model is currently being used to transcribe. "Using X GB of models on disk" at the top of the tab shows your total storage used by downloaded models.

Troubleshooting

CaptionLab doesn't appear in Window → Extensions

Fully quit and reopen Premiere Pro, not just close the window.

The AI engine setup seems frozen

The model download is the longest step (a few GB) — progress is shown in the panel. Don't quit Premiere during setup. If it genuinely hangs with no progress for several minutes, quit Premiere, run `Uninstall CaptionLab.bat`, then reinstall.

"Premiere's scripting engine isn't responding" when clicking Generate

A known Premiere Pro quirk: some workspace layouts don't fully hand scripting access to panels. Try switching to `Window → Workspaces → Editing` and click **Refresh** in the panel, then try again. If that doesn't help, try a different workspace layout. Still stuck? Quit Premiere Pro completely and reopen it.

A model download failed or got stuck

Click **Retry** on that model's card in the Models tab — CaptionLab resumes from whichever files are still missing rather than starting over. If it keeps failing, check your internet connection and try again; some corporate/antivirus networks briefly interfere with large downloads.

⚠ Reinstalling an update?

Always fully quit Premiere Pro *before* reinstalling, not just close the panel. Installing new files while Premiere is still open leaves it unable to run CaptionLab's scripts until Premiere is restarted.

Uninstall

Double-click **Uninstall CaptionLab.bat** (included in the zip). It removes the extension, the background engine, and the downloaded Python/model files.

Support

Email: editlabtools@gmail.com

Website: captionlab.pro

When reporting an issue, please include: Windows version, Premiere Pro version, GPU model (if any), and a screenshot.