

THE "VOICE" GUIDE

VoiceStudio: the free ElevenLabs alternative, set up step by step

You commented "voice", so here's everything from the reel: the official links, what your computer needs, and how to install it and make your first cloned voice. Read page 2 before you download. It saves you from the two most common problems.

GitHub repo	github.com/debpalash/VoiceStudio
Download page	github.com/debpalash/VoiceStudio/releases/latest
Official website	voicestudio.sh

₹0 the free, open-source version (AGPL-3.0)	646 languages in the default voice model, Hindi included	90,000+ GitHub stars (90,419 on 27 Oct 2026)
Local runs on your own computer, no account needed	~10 GB free disk space for the app and its models	v0.5.7 latest release, 26 Oct 2026

What it can do

- ☐ **Voice cloning:** make speech in your own voice from a short recording
- ☐ **Voice design:** describe a voice and get a new one
- ☐ **Video dubbing** into other languages
- ☐ **Transcription** and dictation (speech to text)
- ☐ **Audiobooks and stories** from long scripts

What's in this guide

- 1** Before you start: check your computer
- 2** Download the right file
- 3** Install it (Mac, Windows, Linux)
- 4** First launch
- 5** Clone your voice
- 6** Design, dub and transcribe
- 7** Fixes for common problems

PART 1 Before you start: will it run on my computer?

VoiceStudio runs the AI on **your** machine, so your computer matters. Find yours in this table. It comes from the project's own README.

YOUR COMPUTER	WORKS?	NOTES
Mac with Apple Silicon (M1, M2, M3, M4...)	Yes	Uses the Mac's GPU (Metal) automatically. Needs macOS 13.3 Ventura or newer.
Windows 10 / 11 PC with an NVIDIA GPU	Yes, fastest	Uses CUDA automatically when the NVIDIA driver is installed.
Windows or Linux PC without a dedicated GPU	Yes, slower	Runs on the CPU. The docs say this can be 10-50× slower than a GPU. Needs about 5 GB instead of 9 GB.
Linux x64 desktop	Yes	Snap, AppImage or .deb. No ARM64 Linux packages.
Intel Mac	No	Only the app's interface opens. It can't generate voices locally.
Windows on ARM (Snapdragon X)	Experimental	CPU only, not fully tested yet.
Android / iPhone	Preview	Early phone apps exist; this guide covers the computer app.

What you need (keep it ready)

- ☐ A computer from the "Yes" rows above
- ☐ **About 10 GB of free disk space**
- ☐ Good Wi-Fi: the first launch downloads several GB
- ☐ A quiet room and a few seconds of clean audio of your voice
- ☐ Optional: a free Hugging Face account (faster downloads and bigger models)

Free vs Pro: what's actually free

The free open-source app has every feature in the reel: cloning, design, dubbing, transcription and audiobooks. No account, no subscription.

There is a paid **VoiceStudio Pro** (\$99/year or \$299 once). It adds a commercial licence for client work, more devices and early features. You don't need it to follow this guide.

Read this before using it for work or clients

- The app is free for commercial use, but the **default voice model's weights are licensed CC-BY-NC (non-commercial)**, according to the project's licence notice. Each model in the app's Model Manager shows its own terms. Check them before using the output in paid work, ads or a monetised channel.
- Only clone voices you have permission to use:** your own, or someone who agreed. The project asks for this too.
- It's young software (first released April 2026). Expect the occasional bug, and keep the app updated.

PART 2 Download the right file

1 Open the official Releases page

Go to github.com/debpalash/VoiceStudio/releases/latest on your computer. Scroll to the **Download** table and click the installer for your system.

Download	
Platform	Installer
Windows x64	Installer
Windows ARM64 (experimental)	Installer
macOS Apple Silicon	DMG
macOS Intel	DMG
Linux x64	ApplImage · deb · Snap Store (<code>sudo snap install voicestudio</code>)
Android	APK · AAB
iPhone	Sideload IPA — sign with your own Apple ID

The Download table on the v0.5.7 release page (27 October 2026). Pick your row.

2 Pick your file

Apple Silicon Mac	macOS Apple Silicon → DMG (file ends in <code>mac-arm64.dmg</code>)
Windows PC	Windows x64 → Installer (file ends in <code>win-x64.exe</code>)
Linux	Linux x64 → ApplImage (any distro) or deb (Ubuntu/Debian), or the Snap Store

Mac users: not sure which chip you have? Click the Apple menu → **About This Mac** . "Apple M..." means Apple Silicon. Don't download the "macOS Intel" file on an M-series Mac. It won't work properly.

3 Or use the one-line installer (optional)

If you're comfortable with Terminal, the official script picks the right file, checks it and installs it:

```
Mac / Linux (Terminal):  
curl -fsSL https://voicestudio.sh/install | sh  
  
Windows (PowerShell, not as admin):  
irm https://voicestudio.sh/install.ps1 | iex
```

Only download from the official page

Reels like mine bring fake "VoiceStudio" downloads. Use only the GitHub link above or voicestudio.sh. Each release also lists a `SHA256SUMS.txt` file if you want to verify your download.

PART 3 Install it

Mac (Apple Silicon)

1. Double-click the downloaded `.dmg`.
2. Drag **VoiceStudio** into **Applications**.
3. Open it from Applications.
4. If macOS says it "cannot be opened because the developer cannot be verified": **right-click** the app → **Open** → **Open**.
5. On macOS 15 Sequoia: double-click once, then go to **System Settings** → **Privacy & Security**, scroll down and click **Open Anyway**. You only do this once.

Windows

1. Double-click the downloaded `.exe`.
 2. Follow the setup wizard. It installs just for you, so no admin rights are needed, and you can choose the folder.
 3. Open VoiceStudio from the Start menu.
- No NVIDIA GPU? That's fine. The app detects this and installs the smaller CPU version.

Linux

Snap:

```
sudo snap install voicestudio
```

AppImage:

```
chmod +x VoiceStudio-*  
*-linux-x64.AppImage  
./VoiceStudio-*  
*-linux-x64.AppImage
```

.deb:

```
sudo apt install  
./VoiceStudio-*  
*-linux-x64.deb
```

PART 4 First launch: let it set itself up

1 Wait for the setup screen to finish

On the first launch, VoiceStudio downloads its own Python environment, FFmpeg and AI models. You don't install anything yourself. This is the big download (several GB), so stay on good Wi-Fi and keep the app open until it says it's ready.


2 Optional: add a free Hugging Face token

The models come from Hugging Face. A free token makes downloads faster and unlocks some larger models. Make a free account at huggingface.co, create a token, then paste it in VoiceStudio under **Settings** → **Credentials** and click **Test now**.

3

You land on the Launchpad

This is the home screen. Every tool from the reel starts here: Voice Clone, Voice Design, Video Dubbing, Stories, Audiobook, Voice Gallery and Transcripts.



The Launchpad, from the project's own docs. Your version may look slightly different.

Tip: the Launchpad has a "Try the demo voice" line. Click it once to check everything works before you record anything.

PART 5 Clone your own voice

1

Open Voice cloning

From the Launchpad, click **Voice Clone** (or **Voice** → **Clone** in the sidebar).

2

Give it your voice


Choose a saved voice, or add a **clean reference recording**: drop in an audio file or record one. For the best result:

- record in a quiet room, close to the mic, with no music in the background
- only one person speaking
- speak naturally, the way you want the clone to sound

3

Type your script and generate

Type what you want it to say in the **Script** box, set the language (or leave it on **Auto**), then click **Synthesize audio**, or press **Ctrl/Cmd + Enter**. If the app asks to install a model the first time, say yes.



The Voice cloning workspace: voice at the top, script in the middle, Synthesize audio at the bottom right. From the project's docs.


Making it sound better

A cleaner reference clip helps more than anything else. If it sounds off, re-record in a quieter spot. Hindi works with the default model, and quality varies between languages, so try a few lines first.

PART 6 Design a voice, dub a video, transcribe

Voice design

Open **Voice Design**. Describe the voice you want (for example "warm, middle-aged, Indian English narrator"), or start from a preset such as Narrator or News Anchor. Then type a script and click **Synthesize audio**. No recording needed.




Voice design with ready-made starting points.

Transcription

Open **Transcripts** to turn audio or video into editable text. The app uses Whisper models locally by default.

Video dubbing

Open **Video Dubbing**. Click **Choose file...** (MP4, MOV, MKV, WEBM, MP3, WAV and more) or paste a YouTube or video URL. Pick the language to dub into under **Dub into**, then generate the dub.




The Video Dubbing Studio: original on the left, dubbed version on the right.

Audiobooks and stories

Open **Audiobook** to turn a long script or EPUB into a chaptered audiobook, or **Stories** for scripts with several characters and voices.

SPEED Make it faster

Go to **Settings** → **Models**. The model packs let you trade speed for quality: **Fast**, **Balanced**, **Quality**, **Max**. On a laptop without a GPU, start with Fast. On an NVIDIA GPU or an Apple Silicon Mac, try Quality or Max.



Settings → Models: pick a pack, then install it in one step.

PART 7 Fixes for common problems

PROBLEM	WHAT TO DO
"VoiceStudio cannot be opened" (Mac)	Right-click the app → Open → Open. On Sequoia: System Settings → Privacy & Security → Open Anyway. If it says the app is "damaged", re-download the latest release.
"Not supported on this Mac" or asks for Rosetta	You installed the Intel file on an M-series Mac. Download the macOS Apple Silicon DMG and install it over the old one. Your voices are kept.
I have an Intel Mac	Voices can't be generated locally on Intel Macs. Use an Apple Silicon Mac, a Windows PC or Linux instead.
Generating is very slow	You're probably running on the CPU. Settings → Performance & Device shows whether a GPU is active. On a CPU-only laptop, use the Fast pack, a light voice engine (KittenTTS, Supertonic-3 or PocketTTS) and a small Whisper model.
Model downloads are slow or stuck	Stay on good Wi-Fi and add a free Hugging Face token under Settings → Credentials. Downloads resume if the connection drops.
Not enough space on C: (Windows)	You need about 10 GB free (about 5 GB on a PC without an NVIDIA GPU). Free up space, or move the app's data in Settings → Storage .
The app can't hear my mic (Linux Snap)	Run <code>sudo snap connect voicestudio:audio-record</code> once.
Something else	Search or ask in the repo's Issues and Discussions tabs on GitHub, or the project's Discord (linked in the README).

That's it. Go make something with your own voice.

If this guide saved you an ElevenLabs bill, the download page has an optional way to support my work. Totally optional, and thank you for watching either way.

See you in the next reel. Follow @anothermodernhuman for more free AI tools.

Sources, checked 27 October 2026: github.com/debpalash/VoiceStudio (README, LICENSE-NOTICE.md, docs/install/macOS.md, windows.md, linux.md, docs/languages.md, docs/performance.md, docs/setup/huggingface-token.md), the v0.5.7 release page, voicestudio.sh/pro. App screenshots are from the project's docs folder. Not affiliated with VoiceStudio or ElevenLabs. Features, prices and licences can change, so check the official pages.