

FF | EAI

EAI Create Studio

User Guide

One video teaches a robot a new move

Build: 2026-09-18

Applies to: developer.ff.com/create-studio

Language: English

Contents

- ▶ 1. What this tool does
- ▶ 2. The interface at a glance
- ▶ 3. Preparing your video
- ▶ 4. Step by step
- ▶ 5. Reading the status and error messages
- ▶ 6. Pre-flight checklist
- ▶ 7. My works — and one important limitation
- ▶ 8. Glossary
- ▶ 9. FAQ

1. What this tool does

Upload an ordinary video — a phone clip is fine. Create Studio reconstructs the 3D human pose in it, retargets that motion onto a robot, checks it against the robot's physical limits, and renders a preview. Once you confirm the preview looks right, the motion goes into reinforcement-learning training, and you end up with a downloadable **action package**: a control policy, the audio, and a deployment guide.

The whole path in one line:


Upload video → Configure → Generate (reconstruct, retarget, precheck, preview) → **Your review** → Send to training → Download package

The interface shows this as a four-step bar: **Upload** → **Configure** → **Generate** → **Preview**. When everything is finished, the last step turns green and reads **Done**.

2. The interface at a glance

Open <https://developer.ff.com/create-studio/> . The top bar has three things worth knowing:

Where	Item	What it does
Left	Motion Skills	Featured example actions, filterable by category and searchable
Left	Create	Upload a video and generate your own action
Top right	EN / 中	Language switch; your choice is remembered
Far right	Round avatar	Opens My works



Motion Skills — featured examples. Categories: All / Dance / Sports / Combat / Gestures / Cinematic, sorted by Trending.

Everything on this page is an official example. Clicking a card only shows a "Featured example action" toast — **your own generated actions never appear here**; they live under My works.

3. Preparing your video

3.1 Hard requirements (stated on the upload panel)

- Format: **MP4 / MOV / WebM**
- Size: **max 1000 MB**
- Length: **max 5 minutes**
- Framing: **best with one full-body person in frame**
- Lighting: **good lighting, minimal occlusion**

3.2 What actually raises your success rate

The compliance check rejects motions the robot cannot learn or would fall doing. The two rejections you will actually hit are **too much spinning/turning** and **too much squatting or vertical bobbing**. So:


- **Shoot from the front and keep turns small.** Spins and repeated body rotations get rejected.
- **Keep it vertically steady.** Deep squats and jumps don't survive; airborne moves can't be learned at all.
- Keep the whole body inside the frame for the whole clip.
- One person only. A second body in frame confuses the pose reconstruction.
- Avoid heavy backlight and anything that occludes the performer.
- If you want music, **it has to be in the video** — the audio track is pulled from your upload and synced to the motion.

The generated robot action **matches the length of your source video**; the rhythm is preserved.

4. Step by step

Step 1 — Upload the motion video

1. Click **Create** in the top bar, or the **Create robot action** button on the Motion Skills banner.
2. **Drag your file** into the dashed area, or click **Choose file**.
3. Below the upload area, "How it works · human video → robot action" plays two 4-up examples. Each is a synchronized **source video / sampling / retarget / real robot** view — worth ten seconds of your time.




Step 1 — upload panel with the requirements, and the 4-up "How it works" examples below.

Step 2 — Configure

Picking a file takes you straight to the configuration step: your video on the left, three settings on the right.

Setting	Options	Default	What it controls
Robot model (required)	FF Master / FF Futurist	FF Master	Which robot's skeleton and joint limits the motion is retargeted onto
Motion intensity	Conservative / Standard / Expressive	Standard	How much of the original amplitude is kept. Bigger motion is harder to stabilize
Smoothing	Low / Medium / High	Medium	Jitter suppression. Raise it for shaky footage or partial occlusion

Then click **Generate robot action**. The button first reads "Uploading..." — for a large file that takes a while, which is normal.




Step 2 — source video on the left, robot model, motion intensity and smoothing on the right.

Step 3 — Generate

The generate page shows a live checklist of seven stages, a percentage, four reference clips, and rotating explainer cards while you wait.

#	Stage	What happens
1	Motion captured	AI reconstructs 3D human pose and orientation from your monocular video
2	Retarget to robot	The motion is mapped onto robot joints — human and robot proportions differ
3	Compliance check	Turn speed and squat depth are verified against the robot's limits
4	Preview & audio sync	The robot preview is rendered and audio from your video is synced to it
5	Your review	You confirm it looks right before the expensive RL training
6	RL training	Thousands of trial-and-error rollouts in physics sim learn a stable controller
7	Package	The control policy, audio and deploy bundle are exported for download



Step 3 — stage checklist, progress, and the four reference clips lighting up as the pipeline advances.

Reaching 100% means the automatic part is done and it is your turn to review — it does not mean the action is finished.


On this page you can wait (status refreshes every two seconds) or click **Cancel generation** to start over. If more than **12 minutes** pass without reaching review, the page warns that it is taking longer than usual.

Can you close the tab while it runs? Not advisable. In the current build, My works is tracked **per browser** — see section 7. If you must step away, don't clear your browser data and don't switch browsers.

Step 4 — Preview and review

The preview page puts three panels side by side:

Panel	Content	Note
Source video	Your original upload	Marked "Original · author only" — nobody else can open it
Motion analysis	Reconstruction and sampling result	Check that the pose tracks the performer
Robot preview	The robot in 3D	Interactive, with switchable skins; drag to orbit the camera




Step 4 — the three panels, with "Continue · send to training" as the single primary action.

This is the one place in the pipeline where your judgment is required. Confirm the robot is doing the motion you wanted before continuing, because training is measured in hours. If it isn't right, re-uploading now is far cheaper.

Step 5 — Send to training

Click the primary button, **Continue · send to training**.

- It takes effect immediately. There is **no confirmation dialog** — the button is the confirmation — and it disables itself to prevent double submission.
- A blue status bar appears: **Queued for training · #N · starts in ~X min.**
- The bar refreshes every three seconds and walks through: queued (position, how many ahead, ETA) → training (current stage, time left) → **✓ Training complete — action ready.**
- While that runs, the primary button reads "Queued for training..." or "Training..." and stays disabled. That's expected.




After approval — queue position and ETA; the primary button is disabled until training finishes.

Why training is slow: the robot has to fail its way through an enormous number of simulated attempts before "finish this without falling" becomes reliable. That scale of computation usually takes hours.

Step 6 — Download the action package

When training finishes, the primary button becomes **Download package** and hands you a `.zip` containing:

- **the control policy (`policy.onnx`)** — the trained neural network that computes joint commands live on the robot;
- **the audio** — the track extracted from your source video;
- **DEPLOY.md** — the deployment guide.



Training complete — the primary button turns into Download package.


⚠ The package is not one-tap-on-robot. It is a control policy, a reference motion and a deploy guide. After downloading, an engineer deploys it on the robot with the companion tool and **validates it on a hoist** before the robot performs it on the ground.

The dialog also offers "Download again", so the same action can be fetched more than once.

5. Reading the status and error messages

Everything is reported inside the page; no system dialogs are used. Match what you see against this table:

Message on screen	What it means	What to do
Rejected by precheck · too much spinning / turning	Angular velocity exceeds what the robot can track	Re-shoot front-facing with fewer turns
Rejected by precheck · too much squatting / bobbing	Vertical motion exceeds the limit	Use a clip with less crouching, steadier vertical motion
Rejected by precheck · both	Rotation and vertical motion both out of range	Use a front-facing, low-turn, steady clip
Generation failed	An error during processing; the reason is shown	Follow the message; if the same clip keeps failing, use another
Connection lost	Job status could not be fetched three times in a row	Click "Start over" and re-upload
Upload failed: ...	The file never reached the server	Check size (≤ 1000 MB), format (MP4/MOV/WebM) and your network
 Taking longer than usual	More than 12 minutes without reaching review	Normal for long or complex clips; if it stalls, go back and retry
Primary button greyed, "Queued / Training..."	Queued or training in progress	Just wait — the status bar gives the ETA
Not ready yet on download	Training hasn't finished	Wait for "Training complete"



A rejected job: the icon turns red, the spinner stops, and the button becomes "Start over".

A red icon means this attempt has stopped — there is nothing left to wait for.


6. Pre-flight checklist

Run through this before you hit Generate; it prevents most of the failed attempts:

- ☐ Under 5 minutes, under 1000 MB, MP4 / MOV / WebM
- ☐ Exactly one person, full body in frame, never leaving the frame
- ☐ Mostly front-facing, no continuous spinning or body rotation
- ☐ No deep squats, jumps, rolls or any airborne phase
- ☐ Good light, no occlusion, not shot into a bright background
- ☐ If you want music, the track is in this video
- ☐ The right robot model is selected (FF Master or FF Futurist)

7. My works — and one important limitation

The avatar in the top-right corner opens **My works**, listing the actions you have generated with their current status.



My works — empty until you generate something. Private; these never appear in Motion Skills.

⚠ In the current build, ownership is tracked by browser. The page says so itself: "Actions you generated from this browser (private, not in Explore)". In practice:

- **A different computer, a different browser, or a private window will not show your earlier work.**
- **Clearing browser data loses the list.**
- The work is private and never appears on the public Motion Skills page.

So: **download the package as soon as training finishes** and keep it locally. Don't treat the browser as storage. If you need to hand work to a colleague, hand over the downloaded `.zip`.

8. Glossary

Term	Meaning
Retarget	Mapping human motion onto robot proportions while keeping feet off the floor-clip, hands out of the body, and joints within limits
Compliance check / precheck	The feasibility screen before training, which blocks motion beyond the robot's limits
RL training	Repeated trial and error in physics simulation until the robot learns a controller that stays upright through the motion
Control policy (policy.onnx)	The training output: a neural network that issues joint commands on the robot in real time
Action package (bundle)	The downloadable <code>.zip</code> : control policy + audio + deployment guide
4-up	The way examples are shown: source video / sampling / retarget / real robot, played in sync

9. FAQ

How long does the whole thing take? The automatic part — up to "Your review" — usually finishes within about fifteen minutes, and the page warns you if it passes 12. Training is separate and typically runs for hours; the queue position and ETA in the status bar are the number to trust.

Can I cancel? During generation, yes — "Cancel generation". Once a job has been sent to training the button is disabled; let it finish.

Why can't anyone else see my source video? It is marked "Original · author only" and is visible only to you on your own preview page.

Do I have to choose FF Master? Robot model is required, and the choice is FF Master or FF Futurist. Proportions and joint limits differ between them, so pick the robot you actually intend to deploy on — the retargeted result is not interchangeable.

Can I hand the downloaded package straight to a robot? No. It needs an engineer to deploy it on the robot and validate on a hoist first.

This guide describes the build live at <https://developer.ff.com/create-studio/> on 2026-09-18. Screenshots were taken from that build. The interface is under active development; where the page and this guide disagree, the page is right.