

# Avaturn Releases AVTR-1: First Open-Weights AI Avatar Duplex Model

*Avaturn, the AI lab behind Avaturn.me, releases new open source AI avatar model. beats SOTA models on key benchmarks.*

NEW YORK, NY, UNITED STATES, May 26, 2026 /EINPresswire.com/ -- Avaturn today released AVTR-1, a real-time AI avatar model that sets a new state of the art on key public benchmarks.

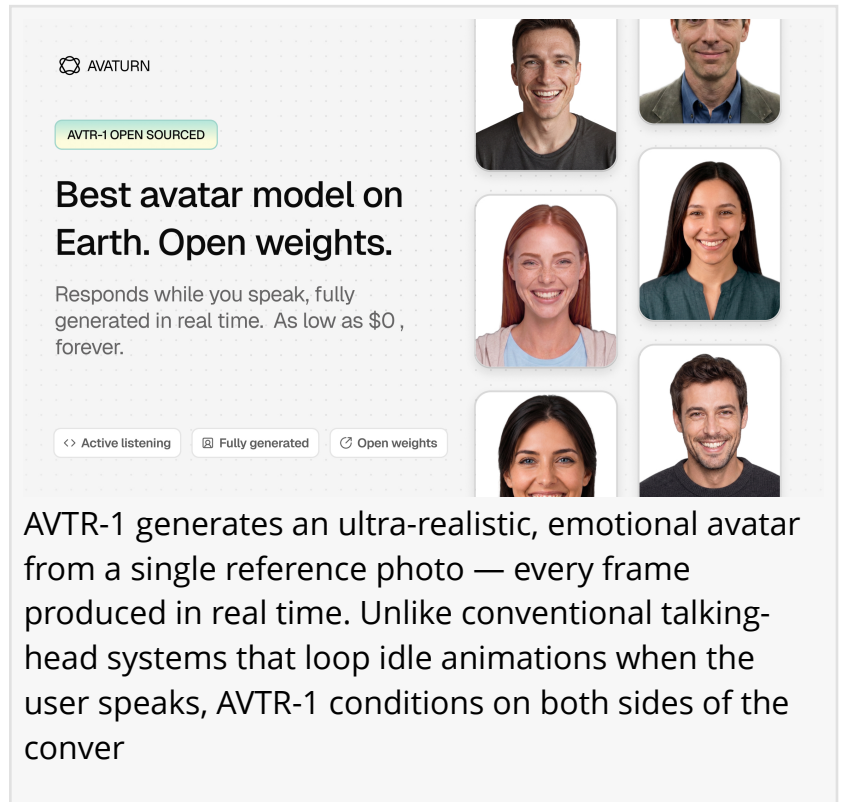
Avaturn is releasing it open-weights.

The new avatar model, the inference stack to run it, and the evaluation methodology are public today on [GitHub](#) and [Hugging Face](#). Weights are released under a community license — free for personal, research, and any commercial use under \$10M in annual revenue. Commercial licensing is available above that, through Avaturn.



Real-time generated video is the next frontier. Every previous wave: text, real-time audio, produced an open layer the category built on. We're shipping that layer today, model and orchestration both"

*Sergei Sherman*



AVATURN

AVTR-1 OPEN SOURCED

**Best avatar model on Earth. Open weights.**

Responds while you speak, fully generated in real time. As low as \$0, forever.

<> Active listening   📷 Fully generated   🔄 Open weights

AVTR-1 generates an ultra-realistic, emotional avatar from a single reference photo — every frame produced in real time. Unlike conventional talking-head systems that loop idle animations when the user speaks, AVTR-1 conditions on both sides of the conver

## Built for interactivity

AVTR-1 was trained on a new architecture, purpose-built for interactivity rather than retrofitted from a general-purpose video model. The narrow focus is what lets the model learn from a fraction of the data general-purpose video models require.

The category of real time interactivity just got its moment. Recent duplex-model announcements from major AI labs have put the framing in headlines: real conversation isn't built on alternating turns — it's built on listening and

responding in the same beat. The same idea applied to video is what real-time avatar research has been chasing for the last year. None have published the a full model that operates with open weights — until today.

"For three years, real-time avatars have been pre-recorded video with a generated mouth pasted on top," said Sergei Sherman, CEO of Avaturn, "We threw out the recording. Every frame is generated, all the face is generated every time . and the avatar is reacting during the entire conversation, not just waiting its turn."

Two things make AVTR-1 different from real-time avatars shipping today.

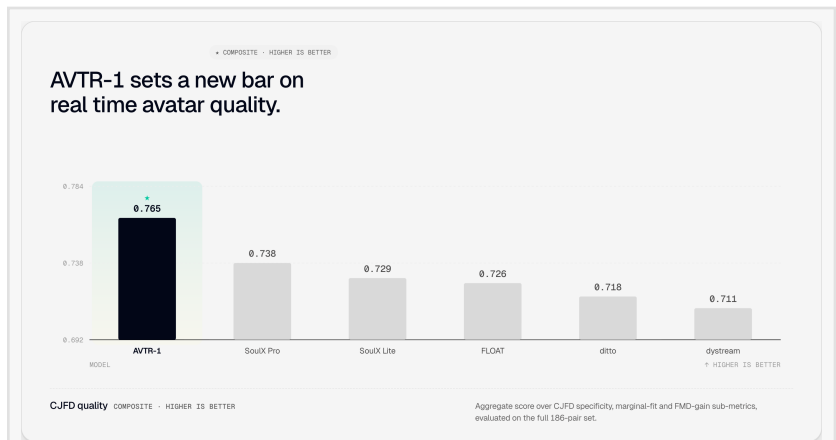
First, the whole face is generated. Not just the lips swapped onto a pre-recorded clip . every pixel of the avatar's face, top of the head to the chin, is generated in real time, frame by frame.

Second, the model is generating all the time, whether the avatar is speaking or listening. . Just like a human on a call, the avatar's face responds to your words and your tone in real time. The brow lifts at word three because you sounded surprised at word three, not after the sentence ends.

Alongside the model, Avaturn is open-sourcing a new real-time streamer, the engine that runs AVTR-1 in real time and accepts any other open-weight real-time video model as a drop-in

Sherman is using the launch to call on the rest of the industry to converge on a public, vendor-neutral leaderboard for real-time AI avatars. "The category needs a transparent scoreboard , one the ecosystem runs together," he said. "Clear, public competition is the only way improvement happens fast. We're inviting every other vendor, every open-source contributor, every researcher to help us build it."

What ships today:




AVTR-1 sets a new bar on real time avatar quality

## Actively Listen, Empathetically Respond

Avatars react in real time — nodding, blinking, smiling while you speak

- Real-time reactions
- Natural nods
- Eye contact
- Engagement cues



AVTR-1 Avaturn real time model is native Duplex with Active listening

Real-time generated video is the next frontier. Every previous wave: text, then real-time audio, produced an open layer the category built on. Avaturn is shipping that layer today: model and orchestration both.

AVTR-1.

A model where every frame is generated in full. Native duplex listening. Sub-200ms end-to-end. Runs on one A100 per session — laptop, data center, cloud. Open weights, training code, architecture paper. Free under \$10M ARR; commercial licensing above. The GitHub repo ships with examples from launch partners Cartesia and Pipecat out of the box.

The Avaturn streamer. The open infrastructure layer for real-time avatars. Any team building text-to-video — any avatar model, any conversation backend, in-house or third-party can plug into the streamer and ship a low-latency experience without rebuilding the orchestration stack from scratch. Plug in your video model on one side. Plug in your conversation backend on the other.

Twelve reference avatars. Model cards, license-cleared, deployable today.

The team behind it

Avaturn is an 18-person AI lab in New York City. The company has raised \$9 million to date, with investors weighted toward notable deep-tech backers. The round was co-led by XTX Ventures — the venture arm of the London-based quantitative trading firm — and Weili Dai and Sehat Sutardja, the husband-and-wife co-founders of semiconductor company Marvell Technology as well as other notable angels.

Avaturn is the team behind Avaturn.me, the world's most-popular 3D avatar platform, used by 9,000+ developers including companies like Disney, Nike, Infosys, and Krafton.

Get the model

GitHub Repo:: (code, inference, evaluation): <https://github.com/avaturn-live/avtr-1>

Hugging Face : (weights, 12 reference avatars): <https://huggingface.co/avaturn-live/avtr-1>

Technical report , [full paper](#) and reproducible benchmarks: <https://avtr-1.avaturn.live/>

Hosted demo: avaturn.live

Sergei Sherman

Avaturn

+1 347-866-1387

[email us here](#)

Visit us on social media:

LinkedIn  
YouTube  
X  
Other

---

This press release can be viewed online at: <https://www.einpresswire.com/article/914997722>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.