


World Labs Atlas: One AI Model That Generates, Reconstructs, and Simulates 3D Worlds

Kabui, Charles

2026-09-08

[Read at ToKnow.ai](#)



The graphic features a dark blue background with a stylized 3D camera rig icon on the right. The main title is in large, bold, white and teal text. Below it, a subtitle in smaller white text describes the system's capabilities. Three key performance indicators are highlighted in separate boxes with vertical bars: '75-94%' in red, '1 min' in orange, and '25.3' in teal. The date 'September 8, 2026' and the 'ToKnow.ai' logo are at the bottom.

World Labs Atlas: One Model That Generates, Reconstructs, and Simulates 3D Worlds

Camera-guided video, 3D scenes from a few photos, and robot simulation in one system

75-94% of human raters preferred Atlas for camera-guided video	1 min of 1440p video generated from a single input photo	25.3 3D reconstruction error, beats specialists (PI3X 28.7)
--	--	---

September 8, 2026 ToKnow.ai

On September 1, World Labs, the startup co-founded by computer vision pioneer Fei-Fei Li, unveiled Atlas, a single model that works with text, images, video, and 3D at once. Each image is pinned to a 3D camera position, so scenes hold together as the viewpoint moves. From one

or a few photos it generates video along any camera path, up to a minute at 1440p, and it rebuilds a real room from as few as two or three ordinary phone photos into an interactive 3D scene, outputting point clouds or [3D Gaussian splats](#). The same model simulates too: it reframes a few phone feeds into “bullet time” shots and generates the image and depth data a robot would see moving through a scanned space.

The pitch is that one generalist beats many specialists. With third-party human raters, Atlas won for following a requested camera move over MiniMax H3 75% of the time, Gemini Omni Flash 81%, and Seedance 2.5 94%. Its sparse-view 3D reconstruction error was 25.3, below the best specialist Pi3X at 28.7. A filmmaker can stage a shot after the shoot, and a robotics team can scan a room with a phone, then generate endless training variations in simulation instead of filming each one. Atlas is in early access for select partners, so these are World Labs’ own numbers until outsiders can reproduce them.

World models are moving from video prediction toward geometry you can move through and hand to robots. One model replacing several specialists would simplify today’s 3D pipelines, if it holds up beyond the demo.

Read More: the story of World Labs’ earlier \$1 billion raise and its Autodesk partnership, [World Labs: Spatial Intelligence AI for 3D World Generation](#).

Sources:

- [Atlas: A World Model for Spatial Intelligence \(World Labs\)](#)
- [Fei-Fei Li’s World Labs debuts Atlas, a world model showcase for advanced spatial intelligence \(SiliconANGLE\)](#)
- [World Labs unveils Atlas, a single AI model that generates, reconstructs, and simulates 3D worlds from just a few photos \(The Decoder\)](#)

Disclaimer: For information only. Accuracy or completeness not guaranteed. Illegal use prohibited. Not professional advice or solicitation. **Read more:** [/terms-of-service](#)