

One picture in, which mesh holds up?

Six game-art pictures went to the image-to-3D models on Scenario twice: once for the best mesh each model can make, and once at fixed triangle budgets. Every mesh was rendered and measured the same way, then put live on one page.



01 · The badger blacksmith, one mesh per model, live on the page

157
meshes on the page

6
pictures

12
models at their best setting

2
budgets: 4K and 10K

Two tests, not one bake-off.

A 2-million-triangle sculpt and a 4,000-triangle game mesh cannot share one ranking. The benchmark asks two separate questions on the same pictures: which model makes the best mesh when triangle count does not matter, and which holds its shape at a fixed budget. Six pictures cover the hard cases: asymmetry and small separate tools, flat panels and lettering, a creature with a shrine and thin flags, an openwork weapon, a full human figure and dense hard-surface detail.



01·Character



03·Vehicle



04·Creature



05·Weapon



07·Anime figure



08·Drill tractor

Test A: high poly

Twelve models, each at its best single-image setting read from its schema: top quality tier, largest texture, PBR on, the highest face count allowed. Judged on geometry, fidelity to the picture and texture. 71 of 72 meshes arrived.

Test B: low poly

Six generators that build low poly directly, each at about 4,000 and 10,000 triangles on five of the pictures, plus a lane that generates a detailed mesh and retopologizes it. Two builder models whose budget field steps by 10,000 ran at 11,000 on the three hard-surface pictures. All 86 meshes arrived.

Reuse what was already at its best.

An earlier comparison (round 2) had already run nine image-to-3D models on six pictures. Round 3 audited its settings against each model's schema before generating anything. Results already at a model's top setting were kept; the rest ran again. Two pictures were swapped for the anime figure and the drill tractor, which every model ran fresh.



08 · The drill tractor in clay, every Test A model, a round-3 picture

Kept from round 2 (22 meshes)

Meshy 7.1, Hitem3D 3.0 at its Quality tier, Sparc3D 2.1, YVO3D at Real 4K texture, Trellis 2 and Pixal3D on pictures 01 to 05. Each tile on the page says which round it comes from.

Run in round 3 (49 meshes)

Hunyuan 3D 3.1 Pro at 1.5M faces, Rodin Gen-2.5 with micro detail and HD texture, and new to the lineup Tripo 3.1, ReconViaGen 0.5, Claude Opus 5.5 3D and GPT-6 Astra 3D. Sparc3D's vehicle ran again because its round-2 result came back black. YVO3D's 8K texture option timed out on five of six pictures, so it stays at 4K.

Every mesh, the same light and angle.

Each GLB went through one Blender 5.2 script, unchanged from round 2. It renders flat silhouettes from 36 azimuths and 3 elevations and keeps the angle whose outline best overlaps the subject in the picture; that overlap is the page's shape-match score. From that angle it renders the mesh textured and as untextured clay, then turns the camera to the side and back, the parts the model had to invent.



The picture



Tripod 3.1· matched angle, clay, back



Measured, not guessed


The same script counts triangles after import (quads count as two), reads every texture and PBR input, and counts open edges after welding the vertices glTF splits along UV seams, so only real holes count.

Clay shows the real geometry

Clay drops every texture, normal maps included. Rivets, lettering and track links that vanish in clay were painted on. At grid size two tiers of one model can look identical; full-size close-ups decide.

Density stops mattering.

Six models stay at 1.4M triangles or more on every picture, yet Rodin at 500K and Claude Opus 5.5 3D at 53K can still win a picture. The differences live in thin parts, lettering and invented backs. More is not always better: Hunyuan at 1.5M faces dropped the prayer flags its 500K run had kept in round 2.



Hitem3D Quality (2M) and Master (5M), clay close-ups

Builders and sculptors


Claude Opus 5.5 3D and GPT-6 Astra 3D deliver editable models made of named parts, often well under their face budget, strongest on hard-surface subjects. The sculpting models keep organic detail closer to the picture.

Lettering and a side test

The 'POST 07' stencil on the vehicle separates the field: some paint it cleanly, two carve a misread 'POST 03' into the geometry. Hitem3D's 5M-face Master tier ran on two pictures beside the grid: identical at grid size, sharper in close-ups.

A budget is a promise to check.

Most low-poly generators land within 10 percent of the target. Tripo 3.1 Smart Low Poly and Tripo Retopology deliver 13 to 88 percent more, likely because the face limit counts quads and the GLB arrives triangulated. The wireframe shows the second difference: some spread triangles evenly, others spend them on the face and hands.



07 · Wireframe at 4K and 10K, twelve of the sixteen low-poly lanes

Generate high, then retopologize


The retopology lane starts from the Tripo 3.1 detailed mesh. Meshy Remesh lands on the budget with an 8K texture but no normal map; Tripo Retopology bakes a normal map from the detailed mesh.

Closed meshes

Rodin Gen-2.5 Fast returns zero open edges on all ten of its meshes. At 11,000 faces and standard effort, the two builder models lose most of the detail they showed in Test A.

Judge it for yourself.

The page puts every mesh of a picture on one shared WebGL canvas, one viewport per tile, starting from the matched angle. Turn them together, switch to clay or wireframe, open any mesh large and read its sheet row. Tiles start as the Blender stills and turn live on demand.



The page: the badger's twelve meshes, live, Test A

Made for the browser

Textures are resized to 1024 px WebP and meshes Draco-compressed. Meshes up to 250K triangles keep their topology; larger ones are simplified to 250K. Hitem3D and Sparc3D give every triangle its own UV island, so their browser copies were decimated in Blender and re-baked from the originals. The sheet always reports the original files.

Take the method

The agent skill packages the renderer and a sheet builder, with the route to run the same two tests on your own picture through Scenario MCP: pick the lineup, price it with dry runs, generate, render and measure.