C-5M Super Galaxy

Human-directed AI agent orchestration

Demetris Dangerfield | NeuroAI Studios. Illustrated case study. October 1, 2026.



Native simulator capture / 0.5.48. C-5M nose, wings and engines on the runway. Historical review capture.

Independent project directed through AI-assisted production. Latest documented owner-test installation: 0.5.74. Complete visual and runtime acceptance remains open.

Project and personal contribution

The C-5M Super Galaxy project is an independent aircraft simulation development effort directed by Demetris Dangerfield under NeuroAI Studios. Its aim is a believable interior and exterior experience in Microsoft Flight Simulator 2024. The project records span reference gathering, geometry reconstruction, texture work, avionics pages, interaction wiring, animations, render infrastructure, SDK packaging, and repeated owner-test installations. This account describes the documented work through October 1, 2026. It does not designate the aircraft as finished or commercially released.

Demetris acted as project owner and director. The chat history records his selection of references and purchased hardware kits, provision of the Windows and DGX Spark resources, direction of Codex and Claude work, simulator observations, and repeated visual corrections. His screenshots supplied concrete acceptance criteria, including instrument-panel slope, pedestal placement, seat scale, missing ceiling coverage, window flicker, and bilateral nose gaps. AI agents authored and reviewed much of the code and asset work. The appropriate professional claim is leadership of an AI-assisted technical production program, with review and integration responsibility, rather than sole manual authorship of every mesh or line of code.

The documented result is an evolving custom aircraft package with substantial implemented cockpit and exterior work, historical native simulator evidence, and reproducible owner-review installations. The current Community manifest was read on October 1 and reports version 0.5.74. Its SHA-256 matches the September 28 delivery record. The latest installation remains an owner-test build; release selection is STOP, and complete native verification and owner acceptance are still open.



Owner photograph of simulator / 0.5.52. Flight-deck overview with yokes, seats, screens and the three-CDU pedestal. Phone colour and screen perspective are retained.

Human-directed AI agent orchestration

Demetris orchestrated the program through a human-owner, coordinator, specialist, reviewer, and integration loop. He set the intended experience, selected authoritative aircraft references, supplied hardware, and inspected test builds in the simulator. His feedback changed priorities and acceptance criteria. Coordinator agents translated those requirements into technical briefs, dispatched specialist lanes, reconciled dependencies, and maintained resumable records. Agent outputs were candidates for him to inspect, correct, or reject.

The September 21 common brief divided owner feedback about the emergency panel, overhead, pilot lighting, Taxi Cam defaults, and EFB into five specialist lanes. Each received a named read-only baseline, a private working folder, a unique prefix for new assets and behavior, a permitted scope, and a candidate-plus-apply-script delivery contract. A specified composition order prevented overlapping changes from being applied arbitrarily. Specialists were instructed to leave the shipping bundle, Community folder, simulator, and other lanes untouched. The coordinator retained integration and installation ownership.

The fire-panel brief went beyond a general prompt. It identified the visible fault in the owner's screenshot, the reference-photo crop, the previous stripe method, the handles and click bindings to preserve, likely causes to investigate, and identical-camera day and night deliverables. Later corrections were batched into bounded rounds. The Taxi Cam brief identified a missing companion process, divided package defaults from persistent auto-start design, and required usable fallback content when the bridge was absent. System-level changes remained with the coordinator and owner.

Demetris could reject a technically green result when the actual nose, rudder, panel shape, scale, lighting, or screen behavior was still wrong. He also controlled model selection and concurrency. The later owner rule capped the team at ten GPT-6 SOL agents across the tools and required one independent judge per changed round. Earlier limits and model assignments were historical policies. The records show bounded team snapshots, not ten agents running throughout the project.

The method evolved through failures. An author could not certify their own candidate. A repaired probe could not immediately certify a geometry change by the same pass. A component result could not accept the complete assembly. The coordinator had to obtain independent observation, exclude stale or incomplete evidence, and route selected outputs into integration. Final simulator and visual acceptance remained Demetris's decision.

Recorded models and changing assignments

Selected Claude transcript metadata records claude-opus-5, claude-opus-5-5, and claude-fable-5. Selected Codex turn metadata records gpt-6-sol and gpt-6-astra. These are identifiers stored in the local project archive. They document model use in the selected sessions; this review is neither a comparative benchmark nor an exhaustive census of every model and subagent. Transcript record counts are not unique agent counts, completion totals, or compute hours.

The task briefs explicitly record background Opus specialists for visual and interaction work, including the fire panel and Taxi Cam. Project rules assign Claude the normal visual-authoring and independent-review lane, and Codex the normal build, release-tooling, and inventory lane. Actual assignments could change with a task. Fable 5 appears in earlier Claude sessions, but this bounded review does not assign it a particular geometry or hardware contribution without a task-level receipt.

The September 27 handoff records Demetris selecting Astra for a focused exterior repair after rejecting the previous attempt. Four Astra agents were assigned nose, rudder, export, and independent review, with root integration separate. Later that day, owner direction switched workers from Astra to SOL and SOL lanes continued the saved work. A selected Codex session records both gpt-6-astra and gpt-6-sol. This illustrates the owner's adaptation of the team; it does not establish that one model generally outperformed another.

The architecture separated frontier-model reasoning from asset compute. Codex and Claude sessions coordinated tools and produced code, plans, and review. Blender, Painter, the MSFS SDK, and the simulator ran on their designated hosts. The DGX Sparks rendered and analysed assets. The reviewed evidence does not show those named frontier models running locally on the Sparks. The Spark GPUs supported the production pipeline while agents directed it through tools.

Hardware, tools, and shared-resource control

The October 1 Windows readback identifies an Intel Core i5-14400F with ten cores and sixteen logical processors, 34,180,726,784 bytes of physical memory, and an NVIDIA GeForce RTX 5080 reporting 16,303 MiB of VRAM. These are present-day inventory readings, not proof of every earlier session's configuration. Windows hosted Blender 4.5, Substance 3D Painter, SDK and package tools, and Microsoft Flight Simulator 2024. It supplied native observations that remote renders could not provide.

The fleet record identifies two NVIDIA DGX Spark GB10 systems, Spark and Spark2, with 121 GB unified memory reported on each and Blender 5.1.0 in the documented render environment. Agents packed scenes on Windows, transferred jobs, used named heavy, light, and CPU queues, and verified returned hashes. One interior GPU job at a time per machine and memory guards constrained work to available resources. Queue and lock discipline mattered alongside hardware capacity.

Substance Painter was another shared resource. A September 22 recovery brief specified an exclusive lock file for Painter RPC, containing the lane and UTC time, and allowed an agent to remove only its own lock. It required checking live Windows processes and both Spark queues before resuming interrupted work. Workers could not kill another lane's job. These controls supported concurrent production without giving every specialist unrestricted control of the desktop or render hosts.

That recovery record describes seven lanes resuming after an API overload. Recovery required reading existing outputs, checking modification times against scripts, verifying returned hashes, and rerunning missing or stale steps. Handoffs, lane briefs, deterministic apply scripts, inventories, and receipts preserved state across sessions and models. This let Demetris redirect work without requiring each replacement agent to reconstruct the project from scratch.

Reference research and aircraft identity

The reference program combined owner-selected cockpit photographs, a Matterport walkthrough of a Lockheed C-5 transport, U.S. Air Force and DoD imagery, DVIDS video, handbook diagrams, exterior three-view references, and selected hardware examples. Captured Matterport views were indexed in a manifest and converted into reference plates. Those plates constrained local part placement and visible structure; perspective images alone did not supply an absolute survey of

every aircraft dimension.

The modern C-5M AMP flight-deck photograph supplied the main layout and colour authority. A photograph of an earlier C-5A/B deck supplied retained-hardware detail, such as the pedestal, yokes, throttle quadrant, seats, placards, and floor. The project explicitly rejected copying its analog main panel or blue panel paint into the modern deck. Keeping the accepted three-CDU pedestal was another owner requirement. A separate night reference called for saturated green legends, amber legends on selected parts, bright displays, and a dim red-orange pilot flood.

The work also researched individual cockpit hardware kits. Project records identify the F16 knobs and switches pack and a guarded toggle switch from CGTrader. These were inputs to private development and adaptation, not proof that a purchased kit already matched a C-5 part. The later rules forbade using an entire C-17 flight deck as the candidate. Historical donor-derived work is credited to Delta Simulations, and the repository records CC BY-NC 4.0 for its original 3D assets. This portfolio includes project imagery and attribution; it does not distribute aircraft source models, purchased kits, or a downloadable aircraft package.

The illustrated public reference examples in this case study are the real-aircraft photograph by Roland Balik and the wheel photograph by Heide Couch, both credited to the U.S. Air Force and identified by their source pages as public domain in the United States. Matterport, copyrighted owner-selected reference photographs, and vendor products are linked as research sources. Their reference plates and purchased source files are retained in the local project archive.

Cockpit geometry and surface production

The cockpit program worked across the main panel, sloped instrument pods and coaming, side consoles, pedestal and throttle quadrant, overhead, pilot and engineer seats, floor, yokes and columns, engineer station, navigator station, windows, ceiling, and upper walls. Additional work covered the corridor, galley, crew lounge, and troop compartment. These were connected spatial problems: moving the floor or seat changed eye position, yoke clearance, pedestal reach, console mounts, and camera framing.

Owner feedback changed the production method. A panel that looked flush or blocky could not be repaired by adding surface detail alone. Later rules required reference plates to be open while modelling, profile-based construction and modifier stacks for hardware, and arrays of a well-made part in place of long hand-typed face lists. World-space vertex measurements, through the object transform, were required because freshly linked object dimensions and bounding boxes had misled earlier passes. Dimensions were labelled measured, derived, or assumed.

A documented pedestal round rebuilt 1,129 nodes at true size, added 30 static nodes, and removed the unwanted grab rail. Those figures describe that source round rather than a fresh census of the current aircraft. The main-panel lanes developed photo-aligned pods and hardware, and the display

lanes investigated the reflection and finish of their glass. Subsequent reviews still found misplaced features, missing placards, surface mottling, incomplete ceiling treatment, and seam defects. The work delivered visible iterations and better construction discipline; it did not establish that every surface reached final visual quality.

Blender 4.5 was the Windows authoring, packing, and diagnostic host. Adobe Substance 3D Painter was used for finish work, including the later bilateral exterior repaint. Tool choice was surface-specific: profiles and procedural geometry suited precise hardware, while Painter suited worn paint, roughness, grime, and similar finishes. The project also used image rectification, texture-channel analysis, UV work, KTX2 preparation, and pixel attribution to separate geometric defects from material defects.



Native simulator capture / 0.5.58. Flight engineer station in the simulator. Historical evidence, not final acceptance.

Displays, interaction, and native observation

The September 19 feasibility record identified eighteen HTML cockpit display surfaces and simulator-variable readers. That was a source assessment of an earlier build, not proof of every display operating in MSFS. Later work extended and refined flight, navigation, engine, systems, CDU, EFB, MADAR, and Taxi Cam pages. The pages used the simulator's HTML and JavaScript gauge path, with real values where sources existed and explicit missing-source treatment where they did not.

The native report for test build 0.5.46 records live main-display painting at KDOV in MSFS 2024. Observed content included attitude and airspeed tapes, navigation values, four-engine data, fuel

and electrical information, and a CDU position and radio page. The same report recorded residual errors: a gear-transit label with the gear down, suspicious idle EGT, selected heading at zero, and intentionally unavailable hydraulic and cabin sources. These are useful observations of implemented pages, not a claim of a complete flight-management computer or study-level system model.

The 0.5.46 control sweep covered 168 Input Events. The report says 167 answered, 146 took and restored a changed value, and 21 momentary buttons returned to their original state; gear testing was skipped on the ground. Side-camera captures showed fire-handle pull and throttle travel. The report explicitly left downstream effects of many overhead controls, Taxi Cam button routing, night behavior, ILS and flight-plan use, and owner mouse clicks unproved. A changing event value is a narrower result than proving the physical aircraft effect.

Further native work documented parked daytime nose-wheel and tail/runway views on both Taxi Cam screens, and a captain-button off/on sequence clearing and restoring an image. Night content remained dark, and continuous refresh, flight cutoff behavior, and the cause of the owner's earlier failure were still unresolved. The display-finish investigation also used native captures to show that daylight glass reflection could wash out an otherwise black page. That made material reflection a concrete correction target rather than treating all washout as a page-colour problem.



Native simulator capture / 0.5.58. Engine and systems pages during a controlled translating scene. This image does not establish a performance improvement.

Exterior, motion, and surface diagnosis

Exterior work addressed wings and pylons, the nose visor and its fixed/moving join, window surrounds, landing gear and wheel finishes, doors and ramps, the tail and rudder, and lighting routes. Native historical captures in this portfolio show the aircraft on a runway and show cockpit equipment through its exterior windows. Diagnostic close-ups show the actual paint and geometry issues that remained. These images are labelled by evidence type and recorded build or candidate.

Nose work demonstrated why an apparently painted line may be a physical form problem. Normal-disabled and paint-only controls, per-pixel node attribution, and fixed/moving source analysis showed that conspicuous outlines and slashes could survive changes to surface maps. Several add-on suppression, bridge, donor, and broad-shell approaches were rejected. The necessary repair had to preserve the real opening joint while correcting the primary fixed and moving surfaces, paint mapping, and motion clearance on both sides.

Rudder work similarly separated isolated component results from the complete assembly. Some receiver and source partitions acquired independently reviewed geometric and continuous-motion evidence. Complete-context renders still revealed crown openings, folds, striations, or intersections. A valid local measurement did not make the entire tail acceptable. The case study therefore treats these as substantial geometry research and partial corrections, with unresolved assembly quality.

The latest 0.5.74 delivery selected a bilateral Substance Painter surface repaint while retaining the prior geometry and animation across four exterior LODs. Painter 12.1.5 and Spark2 Blender 5.1.0 are recorded for that round. Same-camera port and starboard views showed modest reduction of old paint traces, while the prominent upper outline and lower physical slashes remained. The failed new bridge and broad-shell candidates were excluded. The included port comparison is the selected paint-only pair, not the rejected geometry candidate.



Diagnostic Blender render / 0.5.74 paint candidate. Matched port-side paint candidate. Paint traces are reduced, but the mesh-derived outline and slash persist.

Reliable decoding and evidence design

A particularly consequential discovery concerned the MSFS glTF partition contract. StartIndex and BaseVertexIndex are omitted when zero; absence means zero. An earlier reader partitioned only when those keys existed and handed a first primitive the whole shared index buffer. The decoder contract records 499,312 phantom exterior LOD0 triangles and needles up to 28.6 metres. These corrupt diagnostic renders had been mistaken for real aircraft defects. Correct interpretation of the source was a prerequisite to useful modelling and review.

The resulting discipline required a decoder to respect the declared PrimitiveCount, validate index ranges and partition ownership, and demonstrate that the old and corrected readers disagree on the actual asset. A separate crash investigation established the need for single-primitive shipping meshes and the correct carrier hierarchy. The work made export structure and diagnostic fidelity part of production quality rather than assuming a successful file write meant a loadable aircraft.

The project required named gates to be falsified with deliberately broken arms and a healthy control before trusting a green result. Historical failures included a clearance test that did not compare moving partners, UV inspection that dropped zero-area evidence, a bake check that accepted an

inverted map, and a telemetry recorder that discarded valid simulator replies. The final project rules also separated authors from certifiers, and bound evidence to paths, SHA-256 hashes, modification times, and installed versions.

This documentation review did not rerun the historical geometry, export, installation, or runtime gates. Numerical and behavioral findings are attributed to the named archived reports. The local source inventory and asset manifest bind the documents and pictures used here; they do not certify the aircraft. The current manifest readback verifies version identity only. A successful website check or PDF review has no bearing on aircraft release authority.

Distributed rendering and safe delivery

The production infrastructure used two DGX Spark machines as render and analysis hosts, with Windows retained for the SDK, package build, and simulator. The documented fleet used NVIDIA GB10 systems with 121 GB unified memory each and Blender 5.1.0 on the Sparks. Windows-authored Blender packs were transferred with hash checks. Queue scripts separated heavy interior renders, lighter exterior renders, and CPU analysis, using locks and memory guards so work could be resumed without killing another lane's job.

The second Spark setup demonstrated that a copied Blender runtime still needed its shared-library dependencies. It also disproved an assumed direct high-speed fabric link on that machine. Infrastructure decisions were based on the hardware and software actually enumerated. On October 1, both Spark queue status snapshots were idle. This portfolio work did not submit render jobs or alter the render infrastructure.

The aircraft build and owner-review installation route was consolidated behind pipeline/release_cockpit.ps1. Source selection, SDK compilation, compiled correspondence, complete prior-package retention, file verification, and rollback on mismatch were distinct steps. Public release authority remained separately locked. Review installation was authorised for testing and did not imply a public aircraft release or owner acceptance.

The September 28 build receipt for 0.5.74 records SDK exit zero. Its installation transaction records 3,635 verified deployed files and retention of the complete prior 0.5.73 package. The October 1 manifest SHA-256 is 244f1e941c7ff5308e5ff7e1d2f31700981b5003c1899f2ca49c7a2266731964, matching the delivery record. This turn did not rehash all deployed files or rerun the SDK. The subsequent September 28 native attempt obtained no loaded-aircraft bilateral or motion observation because computer-use approval timed out.



Native simulator capture / 0.5.48. Fuselage, wing and engine-pylon view. Historical review capture.

What worked and what remains

The most useful method was to make owner-visible requirements concrete and then narrow the cause with controls. Photo-aligned cameras made shape comparisons useful. World-space measurement improved part placement. Native page and event readback distinguished authored functions from observed ones. Normal-disabled and paint-only views prevented a texture pass from being misrepresented as a geometry repair. Exact file binding made an installation traceable, and retained prior packages made review builds reversible.

The most important unsuccessful pattern was work named as complete before the test actually looked at it. A synthetic control sweep could miss a real simulator-object mismatch; a render could look plausible while clicks, lights, or camera images remained unproved; an isolated clearance result could miss the full moving assembly. The program's response was to preserve the red evidence, broaden the observation, change the method after repeated attempts, and reserve final visual acceptance for the owner.

A September 25 recorder run contained 184 valid samples over 920.11 seconds and was independently accepted for telemetry only. It involved a translating scene with altitude and attitude freeze, early overspeed, and concurrent CPU Blender activity. It does not support a normal-flight or performance-improvement claim. The interior and exterior captures from that run show appearance at their respective times; they do not establish a matched early/late blur comparison.

Open work at this snapshot includes the visible fixed/moving nose join, full rudder assembly quality, current-build native observation, EFB flicker and click behavior, night and camera completeness, intelligible automatic callouts, texture and performance diagnosis, and final owner acceptance. The portfolio's professional value rests in the implemented systems, production engineering, research, iteration, and honest verification boundary. It does not require describing the whole aircraft as finished.

Source coverage and reproducibility

The portfolio review read 34 selected project documents, receipts, and workflow briefs, eight selected Claude root-project sessions, and four selected Codex sessions. The Claude root project contained 47 session files at the inventory snapshot; this was a bounded selection rather than an exhaustive audit of every message and subagent. Project handoff records and exact delivery receipts take precedence over conversational summaries where claims differ. The site uses fourteen curated full-frame images, each bound to its source hash and modification time in the local asset manifest.

Claude's live desktop-code session was checked through its session registry and was working in CFB27Mods. Codex's current chat and other visible project chats were inspected, and both Spark queues were checked. All new authoring, copied records, local previews, PDF outputs, and publication receipts for this task stay under pipeline/runs/c99_portfolio_20261001. No aircraft candidate, PackageSources, compiled aircraft package, Community installation, or shared handoff was edited for the portfolio.

To reproduce the portfolio assets, run prepare_assets.py with the bundled Python runtime. To regenerate the public case-study pages and illustrated PDF, run author_case_study.py in the same lane after the required artifact-operation marker. The local evidence directory contains source snapshots, selected private chat context, an asset manifest, and review records. Private chat context stays outside the website checkout and is excluded from deployment. The final website is a static publication with a detailed project account and an illustrated case study.

Research references

Matterport C-5 walkthrough. Spatial reference for cockpit, stations and cabin. The scan was used as a reference, not an absolute survey. Source.

DVIDS interior B-roll. Jonathon Carnell, 60th Air Mobility Wing, May 5, 2021. Real-aircraft interior and equipment reference. Source.

Real C-5M airframe photograph. Roland Balik / U.S. Air Force, April 16, 2014. Public-domain reference photograph displayed in the gallery. Source.

C-5M wheel photograph. Heide Couch / U.S. Air Force, March 13, 2019. Public-domain wheel reference displayed in the gallery. Source.

C-5 handbook source listing. Project-held technical reference for door, ramp and equipment relationships. Historical C-5 references require variant judgement. Source.

MSFS Blender exporter documentation. Export and simulator integration reference. Source.

Blender 4.5 Bevel modifier. Reference for profile and modifier-based hard-surface construction. Source.

Delta Simulations C-17 project. Credit for historical donor-derived cockpit assets. Repository attribution records original 3D assets under CC BY-NC 4.0. Source.

CC BY-NC 4.0 licence. Noncommercial licence referenced by the project attribution record. Source.

CGTrader F16 knobs and switches pack. Purchased private-development hardware input recorded in the project kit inventory. Source.

CGTrader guarded toggle switch. Purchased private-development hardware input recorded in the project kit inventory. Source.

chickenstrip23 legacy C-5 simulation video. Historical simulation benchmark recorded in the reference pack; not native evidence for this project. Source.