EON AI Ventures

Universal Mirror Direct

Building measured, honest digital twins from photographs alone

How an automated method turns a phone walk-around into a component-level twin whose every part states whether it was measured or guessed — and how it differs from EON's established, expert-guided Mirror Studio.

Table of Contents

1. Executive summary

EON has demonstrated, stage by stage and with measured results on one of its own full-scale industrial training skids, a method that builds a measured, component-level digital twin directly from ordinary photographs. The metric reconstruction stage runs in 27 seconds for 16 photographs on a laptop [5]. The twin's scale was validated three independent ways — against a published ASME pipe standard to +0.3%, against a published strut standard, and against the equipment's own nameplate data [2][7] — and the piping was read off the photographs themselves: 24 runs, each carrying the images that prove it [6].

Two things distinguish the method. First, **automation with a deliberately small human role**: where EON's established Mirror Studio walks an expert through eleven guided steps, Mirror Direct is designed so a person makes three decisions — the rule at the start, revisions by typed comment, and publish [12][13]. Second, **honesty as architecture**: every part is badged with how it is known, everything unestablished is counted rather than dropped, and the twin ships with a square-by-square faithfulness check against the original photographs plus a named list of gaps [1]. EON's August 2026 market scan found no shipped industrial product that goes from photographs to a clean, named, clickable twin [4]. Mirror Direct is a technology demonstration entering productisation on EON's Genesis platform — this paper states plainly, in Section 8, what is proven and what remains open.

2. The problem: three roads that all stop short

Industrial digital twins today are built three ways, and each fails in its own characteristic mode [4]:

- **Classic photogrammetry** (the surveyor). Highest raw accuracy when given 100+ disciplined, well-overlapped photos — but moving highlights on shiny metal still break feature matching, casual capture fails hard rather than gracefully, and the output is a heavy triangle soup with zero semantics: every component must still be modelled by hand [4].
- **Generative image-to-3D** (the guessing artist). Beautiful, complete, game-ready objects — but the hidden side is invented, no model exceeded ~60% surface-detail faithfulness from a single view in the industry benchmark EON's scan recorded (Aug 2026), every output is a normalised mesh with no real-world scale, and the weak spot every comparison lists is exactly complex mechanical assemblies [4].
- **Document-driven twin platforms** (the incumbents). Real semantics — but only from laser scans and engineering documents. Most sites cannot produce a usable P&ID; if the drawing is required, the product is not deliverable [2][4].

EON's market scan of August 2026 — three parallel research sweeps across roughly 55 sources — concluded that the market splits into exactly those camps and that **"photos in → clean, named, clickable twin out" is shipped by no one for industry** [4]. The architecture is academically validated for indoor scenes; the industrial lane is open. EON also ran its own bake-off on its own photographs (one run, one seed, one AI judge — directional, not final): at whole-skid scale every generative model hit the same ~7/10 wall on pipes and valve detail, while at component scale with two close-ups the best reached 9/10 — confirming that scene truth and component detail must come from different instruments [4].

Why it is worth solving: in the accident analysis EON uses across its safety work — 184 fatal industrial accidents — 90% of the crews involved had never seen that plant, 53% were at the equipment for the first time, and 86% of the accidents involved how the work was planned and authorised [11]. Crews should be able to walk the plant before mobilising. That requires a twin per site, per rotation — which is an economics problem: the twin must be buildable from photographs that often already exist, at a cost that repeats.

3. What a twin has to be: components, not models

The single most expensive discovery of the programme looks the smallest. A folder of good 3D models did not add up to a machine. The team's analogy [1][2]: a primitive box is a LEGO brick — plain, but its studs are in known places, so a pipe knows where to land. A scanned or generated model is a pebble — the right shape, beautifully so, with nothing on it saying where the pipe attaches, which way is front, or how big it is in metres. Swap bricks for pebbles and the twin gets better looking and less connectable at the same time.

The rule that fell out, now load-bearing in the build: **a component = model + ports (with axes, in its own frame) + orientation + size**, stored durably — miss any one of the four and the part cannot be seated. And its inversion: **seat by ports, not by bounding box** — place the body so its own port lands on the pipe, instead of placing the body and hoping the pipe reaches [1].

4. The method: twelve stages in four bands

Figure 1 is the pipeline on one page. The stages were not designed on a whiteboard; each exists because the previous one produced something the photographs contradicted [2][3].

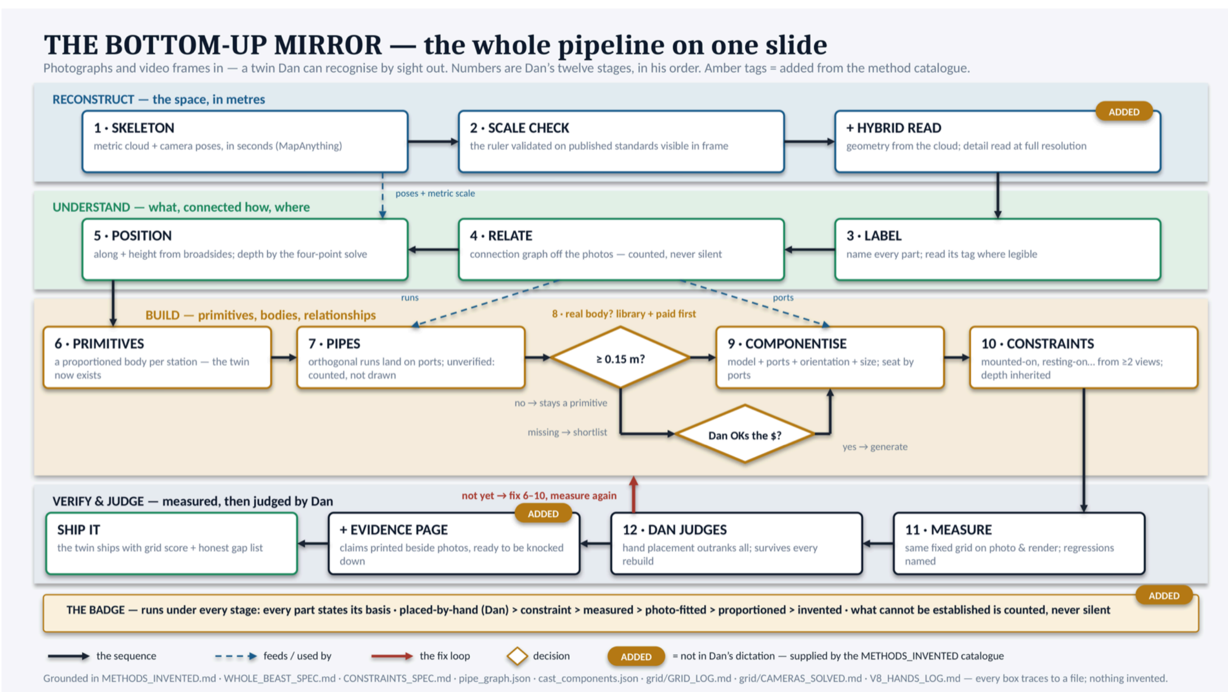

Figure 1 — the twelve stages in four bands. [Internal title says "Bottom-Up Mirror"; regenerate with the confirmed public name before publication.]

Band 1 · RECONSTRUCT — the space, in metres

A feed-forward reconstruction model (MapAnything, Meta — Apache-2.0, self-hosted, so customer photographs never leave the building) turns 16 unposed photographs into a metric point cloud and camera poses in 27.4 seconds on a laptop with no GPU [5]. The scale is then **checked before it is trusted**: any part in frame with published dimensions is a ruler of known length lying inside the photograph. A part made to a standard is better still, because two published numbers check each other inside the same image [1]. Finally, the **hybrid read**: thin pipes ghost at reconstruction resolution (a 42 mm header is four pixels), so geometry comes from the cloud while every small dimension is read in the native full-resolution pixels, foreshortening corrected by the cloud's own geometry — turning a 48-63 mm scatter into a ~1% read [1].

The discovery story behind this band is the paper's cautionary tale. EON's own market deck had prescribed the guardrail two days before the failure: *scale never comes from a generative model or a vision model's eyes — only from the metric skeleton, flange standards, and nameplates* [4]. The build then used something weaker than either: a typed number — a 1.95 m deck length nobody had ever measured. The photographs measured 3.017 m. The fix came from precisely the two sources the guardrail named. The team's own verdict: the broken-then-proven guardrail is the strongest evidence the rule is load-bearing, and it is now a stage in the pipeline rather than a line on a slide [2].

Band 2 · UNDERSTAND — what, connected how, where

Every part is named (with its tag read where legible), and then the **connection graph is read from the photographs**: which port, on which body, connects to which — 24 runs and 45 ports on the demonstration unit, cross-validated across three independently measured photographs to <=0.05 build units [6]. The read corrected the model rather than the reverse: an air riser teed 0.29 m from where the build had it, a “heater return” that turned out to be one object read twice, and four runs with nothing behind them at all [2][6]. Position comes from broadside views for along-the-machine and height; depth comes from a four-point camera solve — with a safeguard that generalises: run the new solver on cameras already known to be good, and if it “corrects” one that was right, throw the result away [1].

Band 3 · BUILD — primitives, bodies, relationships

A proportioned primitive per station makes the twin exist immediately and honestly — a primitive never pretends to know more than it does [2]. Pipes run along the photographed routes and land on ports; runs without both ends established are counted, not drawn [1]. Bodies above the size rule are replaced with real components — library first, then paid generation only on explicit approval — each carrying the full component contract of Section 3. Finally **constraints store the relationship, not the coordinate**: a part bolted to a post has no independent depth — its position is a fact about the post. Nineteen such constraints (mounted-on, in-line-with, resting-on, hangs-from), each citing at least two photographs, are consumed by the build so the fix cannot drift the way a number can [1].

Band 4 · VERIFY & JUDGE — measured, then judged

Photograph and render are cut into the same fixed grid; every machine-bearing cell is judged match / wrong / missing / extra with a one-line note naming its parts. Two rules keep the instrument honest: blank wall and floor are excluded (scoring empty cells would flatter the number), and **the score is never a target** — tune to it and you get fidelity theatre, measured. Its job is to stop regressions: it independently re-found both defects the founder had already caught by eye, which is the test of the instrument [1]. A second rule requires the grid on two axes, because a grid on one view cannot see depth — which pair of views catches a defect identifies the axis the defect lives on [1]. Human judgment outranks everything: a hand placement by someone who knows the machine is the highest evidence basis, and no automatic pass may overwrite it [1]. Every claim is then published on an evidence page beside the photograph it came from, with a how-to-falsify-this line — “so the claims can be knocked down, not admired” [1].

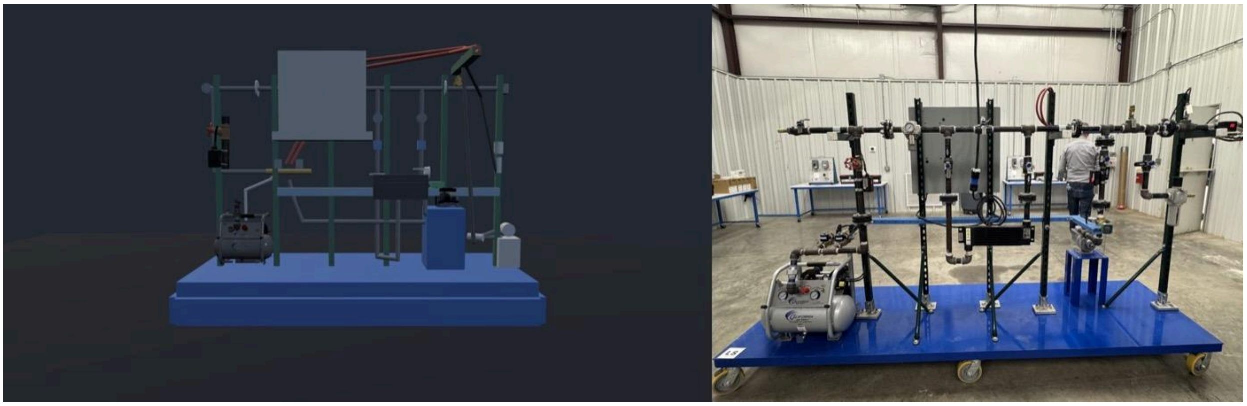

Figure 2 — the demonstration twin (left) against the machine (right), same viewpoint. Real components where evidence supports them; badged primitives where it does not.

Figure 3 — component scale: the compressor cast seated by its ports on the measured frame (left) against the photograph (right).

5. The validation record

Everything below was measured, not estimated. Where a number is a design goal rather than a measurement, it does not appear in this table — it appears in Section 8.

What was checked	Result	Basis
Deck length (the scale that carried the twin)	Typed value 1.95 m, never measured, wrong by ~1.5x. Photographs measured 3.017 m (std 55 mm) across nine independently measured views; 3.000 m adopted.	mapanything/PROGRESS.log; build_cast_ls.py FRAME [5][7]
Independent second route to scale	Pixel measurement against a published strut standard: 3.0 m +/- 0.3. The two routes agree to 0.6%.	Journey deck s10 [2]
Air header outside diameter	Read 42.28 mm vs 42.16 mm published (1-1/4 in NPS, ASME B36.10M): +0.3%.	build_cast_ls.py FRAME [7]; Journey s11 [2]

Strut slot pitch (second published check)	Published 50.80 mm, read 52.26 mm: +2.9%. Raw point map alone: 48-63 mm; the hybrid full-resolution read brings it to ~1%.	Journey s11 [2]; METHODS #5, #6 [1]
Pipe connectivity read from photographs	24 runs: 13 traced end to end, 8 partial with the unseen section declared, 3 counted and deliberately not drawn. 45 ports on named bodies.	pipe_graph.json [6]
Cross-validation of the read	Three photographs measured independently agree to <=0.05 build units on every shared feature.	pipe_graph.json [6]
Independent confirmation of connectivity	The machine's own P&ID, later found photographed on its HMI screen (IMG_5122), confirms the photo-read connectivity station for station.	Journey s8 [2]
Heater position, two methods months apart	Cast lands at 0.147 m; a photo-read primitive built earlier by a different method said 0.146 m — 1 mm apart.	COVERAGE.md [8]
Reconstruction speed	27.4 s inference for 16 photographs on an M1 Pro laptop, no GPU, no cloud upload; install-run-measure in 56 minutes.	mapanything/PROGRESS.log [5]
Camera registration	A 133 mm bias chased for a day was the hand-fitted cameras, not the machine; in the reconstruction, opposed views differ by 3.0 mm.	METHOD_REVIEW.md [9]
Why one scale factor was retired	The twin was wrong by four factors at once: x1.54 along, x1.05 across, x1.36 post height, x1.25 header height — and the same single number hid a main pipe drawn at 26.8 mm against a published 42.16 mm (36% undersized).	build_cast_ls.py v11 header; V11_FRAME_LOG.md [7][10]

6. The honesty architecture

The differentiator is not any single stage — it is that the twin is required to tell the truth about itself, mechanically:

- **Badge everything.** Every part states its basis: measured / photo-fitted / proportioned to a neighbour / placed by hand / invented. The badge outranks the estimate [1].
- **Precedence is enforced in code.** Hand placement > constraint > measured value — a later automatic pass can never silently undo approved work [1].
- **Counted, never silent.** What cannot be established is listed, not dropped: “not drawn: 8 unverified runs — see report” ships with the twin [1][6].
- **The faithfulness check ships with the twin,** as a regression detector, not a grade — together with the named gap list [1].
- **Claims are published falsifiably,** each beside its photograph, with instructions for knocking it down [1].

One retired idea illustrates the culture: the twin was once wrong by four different factors at once (Section 5, last row), so the uniform scale factor was deleted from the code. *The error was not one wrong number — it was the assumption that one number could describe it* [1][2].

7. The human role: Mirror Studio and Mirror Direct

Mirror Studio — the established method — is an expert-guided run of eleven named steps: CAPTURE, FEED, PICK 4, BUILD, INSPECT, REPAIR, TRUE-UP, SPOT, PROVE, MAKE, SHIP. A person approves or deletes every recognised part at INSPECT, and in the MAKE tooling the operator's marks on the photograph are law — the machine may snap them, never override them [12]. It is proven in production use on EON's own training equipment, and it embodies a hard-won rule: real pixels only — never let an AI redraw the part.

Mirror Direct runs the twelve stages of Section 4 automatically. In the productised experience now in design — demonstrated as a fully clickable concept in which every validation number on screen is one the project actually produced — the human makes exactly three decisions, none of them mid-flow [13]:

- **The rule,** at setup: the size threshold above which parts get real bodies, and the generation budget. The build applies the rule and reports what it did — it never stops to ask.
- **Revisions,** by typed comment against the side-by-side photo/twin view — "remove the coil, move the FRL closer to the post" — each applied change entering a ledger badged "changed by comment", plus hand tools for the last mile, with hand placement outranking the machine permanently.
- **Publish,** one button, to EON's Genesis platform — the twin shipping with its faithfulness check, its named gaps, and the badge on every part.

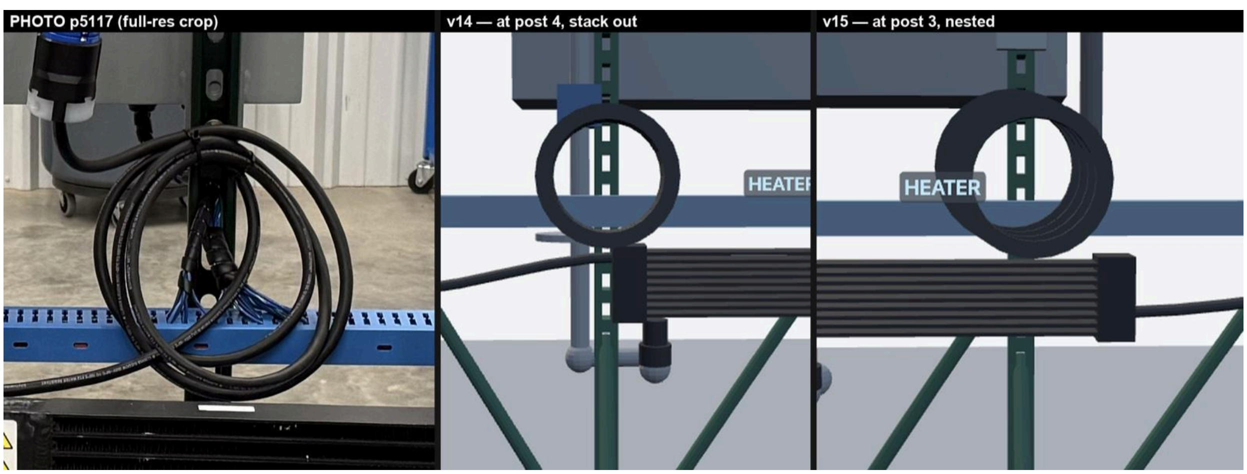

Figure 4 — the revision loop in evidence: the photograph (left), the build with the coil placed one bay off (centre), and the corrected build (right) after the photographs were re-measured. The correction is logged, badged, and survives rebuilds.

The comment loop is deliberate self-portraiture: plain words in, changed twin out, the photographs judge — the exact working loop that built the method itself [13].

8. What remains honestly open

This section is the paper's warranty. Nothing below is claimed as finished [1][2]:

- **End-to-end speed is a design, not a result.** The first build took five weeks of elapsed calendar because it was discovery; the known recipe is designed to run in an afternoon and several stages are measured (reconstruction 27 s; geometry ~1 hour) — but all stages have never been run in one sitting on an unseen machine.

- **Depth was the weakest axis.** Two instruments (a geometric 3D crop and the four-point camera solve) are now proven in use, but two of the four judging cameras still carry hand-set distances — badged in the source, not hidden.
- **One regression is not yet diagnosed.** After the rebuild at true scale, the worst pipe-run gap fell (631 mm to 309 mm) but the median gap on drawn runs grew (2.5 mm to 9.7 mm), and fewer runs close at both ends (11 of 24, against 15).
- **The post count is under challenge** — a re-read says five posts on uneven bays, the twin on disk carries six on an even pitch. Not yet settled, exactly the pattern this paper describes.
- **Six of nineteen stations wear real cast bodies;** the other thirteen are honest primitives standing in the right places.
- **The absolute scale rests on one published dimension** (the 41.28 mm strut section). One tape measure across one post face closes it; nobody has done it.
- **Ports are still hand-defined per part.** Automatic derivation (most promising: from the pipe trace itself) is designed, not demonstrated.
- **Minimal human intervention on a working customer site** is asserted, never demonstrated — access, lighting and time are not ours to choose there.
- **The survey-grade fallback keeps a heavy footprint.** Where a customer needs survey-grade dimensions, classic photogrammetry (RealityScan-class) remains the instrument — with its GPU/Windows compute requirements and disciplined 100+ photo capture [4][14].

9. From demonstration to product

The demonstration twin is public: the showcase at https://eon-sketch-demo.pages.dev/showcase/bottomup/ carries the twin, its photographs, and the badge on every part [15]. Productisation runs through EON's Genesis platform: published twins are hosted there at full resolution (Babylon-based viewer, no showcase size limits), and the Universal-to-Genesis integration is the stated direction [13][15]. Mirror Studio remains the proven path in production use while Mirror Direct completes the open items of Section 8. The two are one product family: the same badges, the same evidence discipline, the same judge — the photographs.

10. Sources and basis

All sources are internal project records on EON's build log unless stated. Numbers in this paper appear with the bracketed reference of the record that carries them.

- [1] METHODS_INVENTED.md — the running method catalogue; every method marked PROVEN / PARTIAL / PROPOSED with its evidence files.
- [2] Bottom_Up_Mirror_The_Journey.pptx (31 Aug 2026) — the research log: the fourteen discoveries and the measured record.
- [3] Bottom_Up_Mirror_Pipeline_Flowchart.png / build_pipeline_flowchart.py — the twelve stages in four bands (Figure 1).

- [4] UNIVERSAL_MIRROR_Market_Analysis_Deck_2026-08-27.pptx — the August 2026 market scan: ~55 sources, three sweeps, five architectures; the guardrail; the bake-off.
- [5] mapanything/PROGRESS.log — 3.017 m (mean of nine independently measured views, std 55 mm); 27.4 s inference for 16 views; 56-minute run, M1 Pro, no CUDA.
- [6] pipe_graph.json / pipe_evidence.html — 24 runs (13 observed / 8 partial / 3 counted-not-drawn), 45 ports, cross-validation <=0.05 build units.
- [7] build_cast_ls.py FRAME dict — the adopted 3.000 m and the typed 1.95 m it replaced; 42.16 mm (ASME B36.10M) and 41.28 mm (strut) published values; every entry carries MEASURED / COUNTED / PUBLISHED.
- [8] COVERAGE.md — heater 0.1472 m measured vs 0.146 m photo-read (the 1 mm agreement).
- [9] METHOD_REVIEW.md — the port-axis requirement; the 133 mm registration bias and its dissolution.
- [10] V11_FRAME_LOG.md — the frame rebuild at true scale; the four anisotropic factors.
- [11] twin-modalities/Twin_Modalities_Analysis_DRAFT.md (26 Aug 2026), citing slide 5 of EON's v18 sales deck — 184 fatal accidents: 90% never saw that plant, 53% first time at the equipment, 86% planning-and-authorisation. [Brita: confirm the underlying dataset is citable publicly.]
- [12] DECK — From one screen to one run.pptx (31 Jul 2026) — the established eleven-step guided run; INSPECT approval; SmartCut rules.
- [13] UX_BOTTOMUP/CONCEPT_v2.md — the three-decision experience, the comment loop, publish to Genesis; every validation number on its screens is one the project produced.
- [14] REALITYSCAN_RUN.md — the survey-grade path's compute requirements.
- [15] DAN_REQUEST_LEDGER.md (1 Sep 2026) — the public showcase live and verified; Genesis hosting direction.