

EON RECOGNITION

How we built a real 3D twin from field photos

Skeleton · Crops · Scale — in plain English, and in the technical detail.

Skeleton

3D shape

Crops

name each part

Scale

real size



WHAT WE STARTED FROM · the real Pump skid, shot on a phone.

STAGE 1 - RECOGNITION

IMAGES
MULTIPLE ANGLES



2D IMAGE CROPS
OF COMPONENTS



SCALE BASED ON
KNOWN COMPONENT



MATCH MAKING


STAGE 2 - CONFIGURATION

SMART 3D LIBRARY

- 51 CLASSES
- 300+ SUB-COMPONENTS
- 8000 PARTS



ASSEMBLED 3-D TWIN



INTELLIGENT PIPE
CONFIGURATION
(AVEVA-LIKE)
AI PIPE BUILDER

- RIGHT DIMENSION
- SNAP FUNCTION
- COLLISION AVOIDANCE

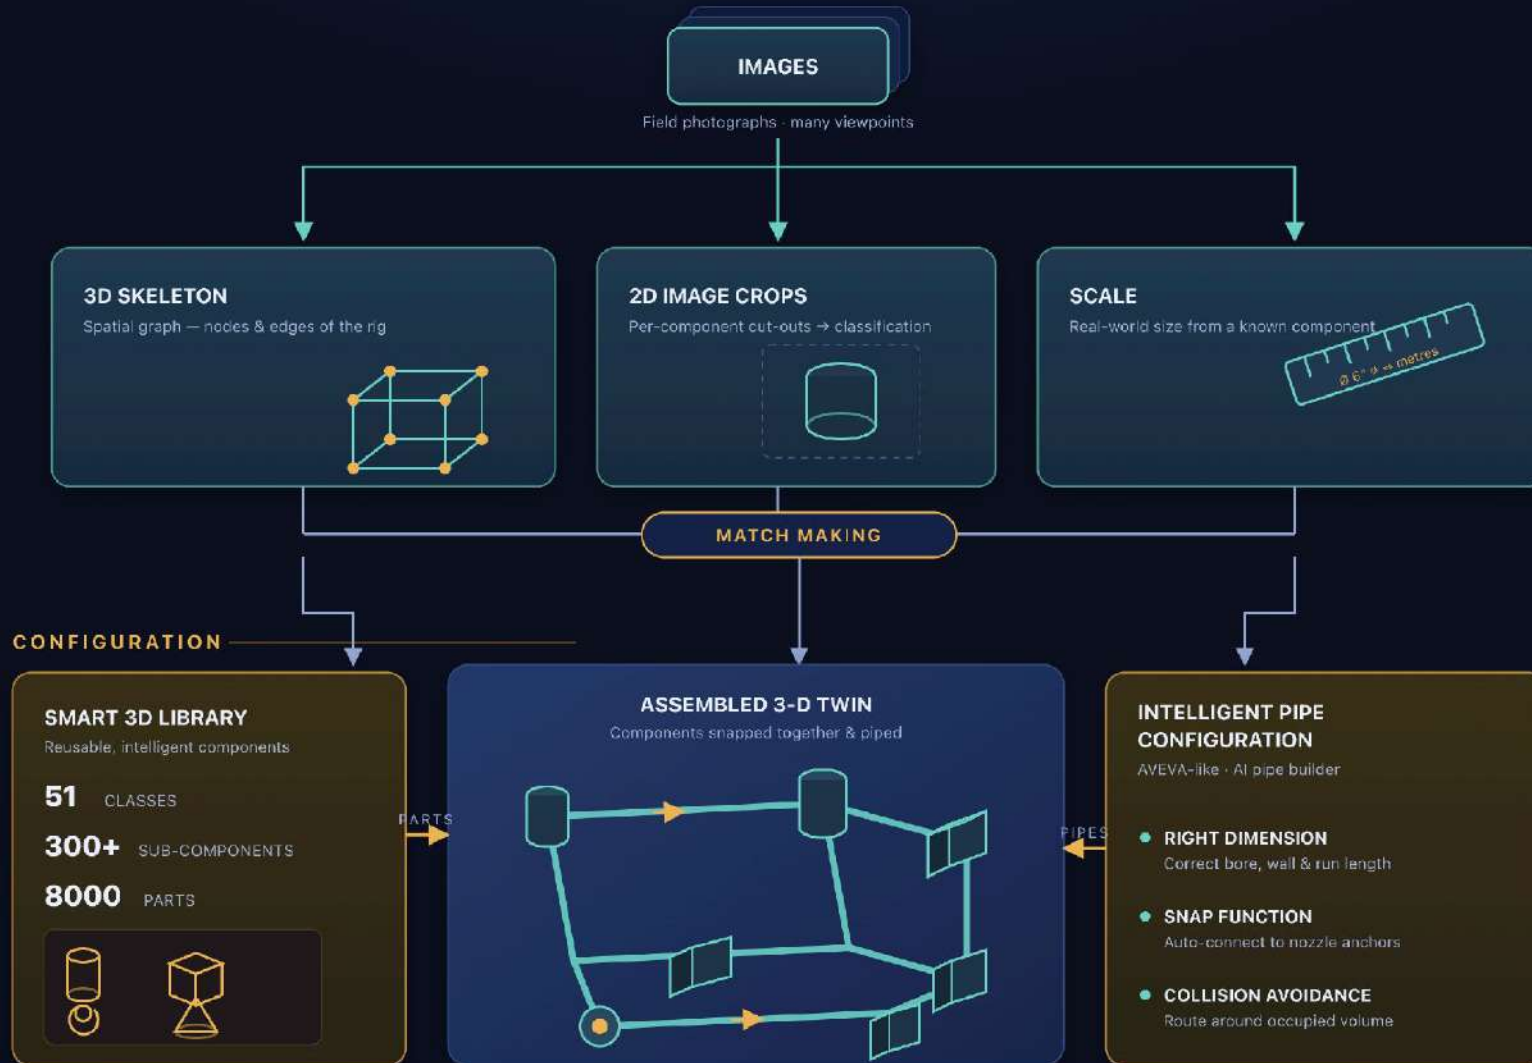
EON UNIVERSAL Photo → Digital Twin · System Architecture

From multiple field photographs to a fully assembled, intelligently-piped 3-D facility twin. Two stages: **Recognition** understands what & where; **Configuration** rebuilds it from the Smart 3D Library and routes the pipework.


■ Recognition stage ■ Configuration stage ■ Data / geometry flow

RECOGNITION

◌ MULTIPLE ANGLES



STAGE 1 - RECOGNITION



MATCH MAKING

STAGE 2 - CONFIGURATION

SMART 3D LIBRARY
Reusable, intelligent components

51 CLASSES
300+ SUB-COMPONENTS
8000 PARTS

PARTS



PIPES

INTELLIGENT PIPE
CONFIGURATION
(AVEVA-LIKE)
AI PIPE BUILDER

- ✓ RIGHT DIMENSION
Correct bore, wall & run length
- ✓ SNAP FUNCTION
Auto-connect to nozzle anchors
- ✓ COLLISION AVOIDANCE
Route around occupied volume

FLOW: > IMAGES (MULTIPLE ANGLES) → RECOGNITION (SKELETON · CROPS · SCALE) → MATCH MAKING → CONFIGURATION (LIBRARY · PIPES) → ASSEMBLED 3-D TWIN

EON UNIVERSAL

Photo → Digital Twin · System Architecture

IN PLAIN ENGLISH

We turned a stack of phone photos into a real 3D twin

A flat drawing tells us what the parts are and how they connect — but not where they sit in space. The photos hold that missing information. Three simple moves unlock it.



1 • A phone photo of the rig



2 • Rebuilt in 3D, checked vs. the photo



Skeleton THE 3D SHAPE

We work out where the camera stood in every photo and rebuild the rig in 3D from the pictures themselves — no guessing off the flat drawing.



Crops NAME EACH PART

We cut a clean close-up of each part, then look at it from many angles at once — like walking around it — and cross-check against the drawing's label.



Scale REAL-WORLD SIZE

A photo model has no ruler. We measure parts whose true size we already know, and use several together so they check each other.

The honest part: *we never trust one source alone. A part is only believed when the photos and the drawing agree.*

The Skeleton — 3D from photos, not the drawing

A P&ID carries connectivity and identity — not geometry. So 3D must be reconstructed from the images.

- Structure-from-motion via COLMAP (pycolmap): SIFT features (~3,700 / photo), exhaustive matching, incremental mapping with global bundle adjustment + EXIF-refined intrinsics.
- Each part's position AND extent is read by back-projecting its 2D box into the point cloud.
- Near-duplicate camera angles (< 5° baseline) are excluded to keep the solve honest.

32 / 54

photos in one 3D frame

3,861


3D points

0.49 px

mean reproj. error



GHOST OVERLAY · 1,423 reconstructed points land on the real frame, pipes, valves & compressor.



POINT CLOUD · top view of the reconstructed rig and the camera path around it.

Layman: a flat drawing can't say WHERE things sit — we rebuild the rig in 3D from the photos, then prove it by projecting the result back onto a real frame.

The Crops — identify each part from many angles

Re-cut from full resolution (~50 px → 400–1800 px), then each uncertain part is checked from ~14 angles by Gemini 2.5-flash against the P&ID tag. A part is “verified” only when photos (≥ 0.85 , ≥ 2 angles) AND the tag agree.



VERIFIED

Pneumatic Hoses — the multi-angle check corrected an earlier wrong “electrical cables” guess.



RECOVERED

Real air compressor behind — recovered from the high-res originals from various angles.



REJECTED

This “part” is actually a person — rejected. (2 such cases caught.)



CONFLICT

Photo shows a hand-wheel valve; tag says solenoid — recorded, not overwritten.

Result: verified parts 5 → 14 · 1 wrong guess corrected · 4 tag-vs-photo conflicts flagged · 2 “parts” were people.

The Scale — real-world size by multi-anchor consensus

Structure-from-motion is scale-free — right shape, no ruler. We recover metric scale from parts whose true size we know, and require several so they cross-check.



Branded → catalog size



Deck length & standard pipe OD

1

Primary — Library ground truth

Our component Library's exact, known dimensions.

2

Free — standard pipe / flange OD

Outer diameter implied by the P&ID line size.

3

Branded — catalog nameplate

Spec of an identified part (e.g. California Air Tools compressor, Schneider Harmony HMI), via Gemini + Google.

4

Optional — one tape measure

A single physical measurement, if available.



The consensus

Each anchor's span implies a scale. **Take the median, report the spread as an error bar, flag outliers instead of averaging. Two independent anchors required.**

Layman: a photo can't tell 2 m from 3 m — so we measure parts we already know the size of, and use several so they check each other.

HOW IT FITS TOGETHER

Division of labor, and the discipline that keeps it honest



Recognition delivers

- the Skeleton (3D frame)
- the Names for each part
- the clean Crops
- the metric Scales



Configurator places

- drops the correctly-scaled Library components into the frame
- verifies each placed component actually looks like its crop



Verification discipline

An adversarial agent reviews every step. And we always inspect the full-resolution image ourselves before trusting — or overriding — any AI verdict.

Bad data in → bad twin out. We fixed the data first — then the twin followed.