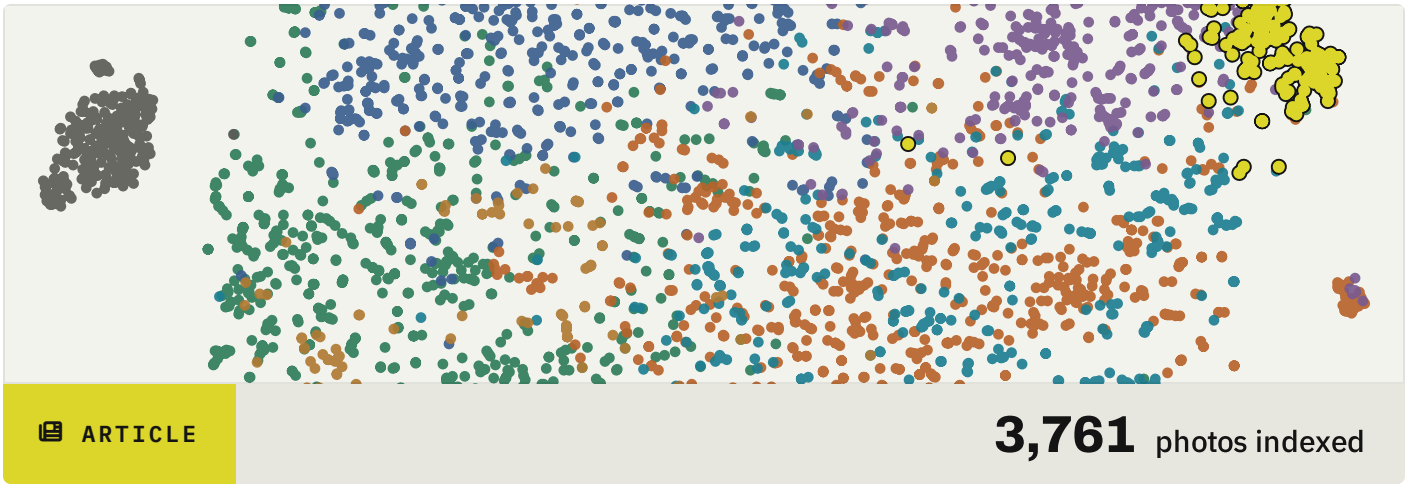
Finding the right photo: 3,761 site photos, searchable by what they show

A site visit produces hundreds of photos. The contamination report needs the right forty. Here is how AVA made its photo archive searchable by content — on its own server — and what we measured.

OCTOBER 6, 2026 · 3 MINUTE READ · BY RICARDO A. CALDERÓN LÓPEZ



The picture: Each dot is one indexed photo, placed by the vision model so that photos that look alike sit together. The colours mark eight groups of similar photos.

The problem: evidence buried in folders

A contamination report for a HUD environmental review ends with a photo log and a list of numbered observations: the tanks, transformers, drums, stained soil and drains that the visit found on and around the site. On a large industrial site one visit can produce more than a thousand photographs — from several phones, sometimes over WhatsApp, often without captions.

Choosing the report-worthy ones used to mean scrolling folder by folder. The photo of the corroded tank shell existed; finding it was the job.

What we built

Photo Index puts every project photo in one searchable index — SharePoint folders, photos from our Visit Checklist phone app and images from approved reports. Each photo is processed on AVA's own server:

- A **vision model (CLIP)** turns every image into a numeric description, so the analyst can search by what a photo *shows* — "above-ground tank", "pad-mounted transformer", "manhole" — or ask for photos that look like this one.
- **OCR** reads the text inside the photo — a stenciled capacity, a warning label, a pole tag — and makes it searchable too.
- Duplicates are recognized by content, and photos are grouped by capture time and GPS position.

3,761	**1,209**	**2,512**
distinct photos from 16 projects	photos with readable text, found by OCR	photos with a GPS position

The same index serves the Image Viewer, the Visit Checklist and the Site Visit Report, so a photo the analyst marks for the report (📋) is marked everywhere.

Teaching the search what a "tank" looks like

Generic image search is good, not perfect. On one project with 133 known tank photos, a plain Spanish query ("tanque") found 20% of them in the top 100 results. An English prompt ensemble reached 32%. A simple model trained on labelled examples taken from the captions of AVA's approved reports reached **54%** — and the analysis puts the ceiling for this collection near 75%.

That is why the analyst stays in charge: the search proposes candidates and the analyst picks.

From photos to a photo log

For the contamination report the analyst reviews contact sheets of candidate photos per topic, groups shots of the same feature, and drafts the entries — caption, description, location — in the register of AVA's approved reports, with numbered observations for the findings section. Uncertain items are written as "not verified" or "undetermined", and dates from the regulatory databases are checked against the EDR report.

In a two-day pilot on a CDBG-MIT warehouse project, **1,276 photos** became a **32-entry, 83-photo log with 19 findings**. A final coverage check over the photos left out found three features the first pass had missed — a cut-open corroded tank shell, a unit discharging to the ground and an aerial of the yard.

> **Where the data goes.** Indexing — the vision model and the OCR — runs on AVA's server; the photos are not sent to an outside service to be searchable. Drafting steps that use a cloud model are chosen case by case, and nothing they write reaches a client without review.

What it changes

Approved AVA reports average about 12 entries and 36 photos. The pilot log was nearly three times as thorough, and every one of its revisions is saved. The work still ends with a professional reading the record — but the time goes into judging the evidence, not hunting for it.

AVA Environmental Consultants, Inc. · PO Box 361300, San Juan, PR 00936-1300 · 787-505-4981 · augusto.velazquez@avaferpo.com · avaferpo.online