

cs-imageindex

Manual

Index photos and videos into a searchable CSV - EXIF, GPS, places, faces, scene descriptions, duplicates

Version 0.4.0 · October 13, 2026

Contents

1. Overview	3
2. CSV columns	3
2.1 Metadata and quality	3
2.2 Faces and scene	3
3. Supported formats	4
4. Usage	4
4.1 Options	4
4.2 Examples	5
5. Face recognition	5
6. Geocoding and threading	5
7. Installation and builds	5
8. Security and privacy	5
9. License	5

1. Overview

cs-imageindex scans a folder of photos and videos (recursively) and writes one CSV row per media file. The index can then be used to build selections - folders of links to matching images - without ever touching or duplicating the originals.

Suggested folder layout:

```
/tank/images
/tank/images/data/*      (originals)
/tank/images/index      (search index, e.g. index.csv)
/tank/images/selections/* (e.g. selections/selection_1)
```

Workflow:

1. Index all images and videos as a napp-it CS job (restartable, skips already indexed files with --resume).
2. Build a selection from the index (a folder of links to the matching images; originals are never modified).

The tool is a single static Rust binary plus a bundled minimal static LGPL ffmpeg/ffprobe. There is no runtime to install and no system library dependency.

2. CSV columns

2.1 Metadata and quality

Column	Meaning
media_type	image / raw / heic / video
date_taken	EXIF DateTimeOriginal
gps_lat / gps_lon	EXIF GPS in decimal degrees (empty if none)
place	reverse-geocoded place name via OpenStreetMap/Nominatim (best effort)
camera	EXIF Make + Model
resolution	pixel dimensions, e.g. 4032x3024
duration	video duration in seconds
orientation	landscape / portrait / square
flash	yes / no / empty
blur_score	Laplacian-variance sharpness (higher = sharper, relative scale)
phash	64-bit perceptual hash (hex) for near-duplicate detection
duplicate_group	shared group number for near-duplicates (Hamming distance within threshold)
is_screenshot	yes if there is no camera Make/Model EXIF (heuristic)

2.2 Faces and scene

Column	Meaning
people	comma-separated names matched against reference photos
unknown_faces	detected but unmatched faces
face_count	total faces detected
tags	scene tags from the vision LLM
ocr_text	text the vision LLM could read in the photo
description	one-sentence scene description

3. Supported formats

- **Still images:** jpg/jpeg, png, bmp, tif/tiff (matched by content).
- **RAW:** cr2, cr3, nef, arw, dng, orf, rw2, raf, pef, srw - decoded in pure Rust, no external tools.
- **HEIC/HEIF:** decoded through the bundled ffmpeg; EXIF is read directly from the file.
- **Video:** mp4, mov, avi, m4v, mkv - one representative frame runs through the same pipeline; creation date, duration and GPS (including the ISO-6709 tag of iPhone .mov files) come from the container.

If the bundled ffmpeg is removed, ffmpeg/ffprobe on PATH is used. If neither exists, HEIC and video files are skipped with a note at startup; everything else runs normally.

4. Usage

```
cs-imageindex --folder /path/to/photos --out index.csv \  
  --refdir /path/to/reference --models-dir /path/to/models \  
  --legacy-cfg /path/to/cs-aihelp.cfg
```

4.1 Options

Option	Meaning
--folder	folder to index (recursive)
--out	output CSV path
--refdir	reference/<Name>/*.jpg for face matching (omit = no recognition)
--models-dir	directory with yunet.onnx and sface.onnx (default: models next to the binary)
--config	own standalone config file (see --print-config-example)
--legacy-cfg	fallback: endpoint2/model2/api_key2 from a napp-it cs-aihelp.cfg
--provider / --endpoint / --model / --api-key	vision provider settings
--ollama / --ollama-model	local Ollama endpoint and model (default llama3.2-vision)
--no-vision	skip scene descriptions (location + faces only)
--no-geocode	no network call to Nominatim
--no-dedup	skip duplicate grouping (saves an $O(n^2)$ pass)
--dedup-threshold N	Hamming threshold 0-64, lower = stricter (default 6)
--threads N	worker threads, 0 = auto (up to 4)
--resume	skip files already in an existing --out CSV and append new rows
--print-config-example	print an example config file

Priority of the vision settings: CLI flags, then the environment variables CS_IMAGEINDEX_PROVIDER / _ENDPOINT / _MODEL / _API_KEY / _MAX_TOKENS, then --config, then --legacy-cfg. If nothing resolves, vision is silently skipped (same as --no-vision).

4.2 Examples

```
# Local Ollama instead of a cloud provider
cs-imageindex --folder ./photos --out index.csv \
  --ollama http://127.0.0.1:11434 --ollama-model llama3.2-vision

# No network calls at all (EXIF + faces only)
cs-imageindex --folder ./photos --out index.csv --no-vision --no-geocode

# Resume an interrupted run
cs-imageindex --folder ./photos --out index.csv --resume

# Stricter duplicate grouping, or none
cs-imageindex --folder ./photos --out index.csv --dedup-threshold 2
cs-imageindex --folder ./photos --out index.csv --no-dedup
```

5. Face recognition

Place one or more reference photos per person in `<refdir>/<Name>/*.jpg`. Detection (YuNet) and recognition (SFace) are ONNX models run through the pure-Rust tract engine - no OpenCV, no onnxruntime. The detection and alignment math is a from-scratch port of OpenCV's upstream C++. Both models are bundled in `models/` from the OpenCV Zoo (MIT and Apache 2.0, license files included).

6. Geocoding and threading

Reverse geocoding uses the public Nominatim API and respects its policy: about one request per second, a descriptive User-Agent and an in-process cache keyed to roughly 1 km. Geocoding is rate-limited globally regardless of the thread count. Disable it with `--no-geocode`.

Media files are processed on a small worker pool (`--threads`). Vision and geocode calls are shared safely; each thread has its own vision API connection.

7. Installation and builds

Release archives follow the napp-it CS convention `<asset>-<platform>.<arch>.tar.gz` (e.g. `cs-imageindex-linux.amd64.tar.gz`) and are installed from napp-it CS under About > Download `cs-tools` into `data/cs_server/tools/cs-imageindex/<platform>.<arch>/`.

GitHub Actions builds Linux, Windows and macOS (arm64 + x86_64) on every `v*` tag. illumos/OmniOS is built by hand with `illumos/cs-imageindex_omnios_1a.sh` (also in `cs-scripts`) and uploaded to the same release; it is fully self-contained. Run the tests with `cargo test`.

8. Security and privacy

- With a cloud provider every photo is sent as a base64 JPEG to the configured endpoint - private images leave the machine. For fully local processing use Ollama (`--ollama`) or `--no-vision`.
- Prefer `--config` or `CS_IMAGEINDEX_API_KEY`; `--api-key` is visible in the process list.
- Directory symlinks are not followed, the scan is depth-capped at 64 levels; file symlinks are indexed as leaf nodes.
- Cell values starting with `=`, `+`, `-` or `@` are written verbatim (the file column is a path). Open the CSV in a text editor or data tool, not by double-click in a spreadsheet.

9. License

BSD 2-Clause, like the other standalone `cs-tools`. FFmpeg is bundled under the LGPL (see `LICENSE-ffmpeg.txt`).