Privacy Keeper

Version 3.0 — User Manual

May 27, 2026

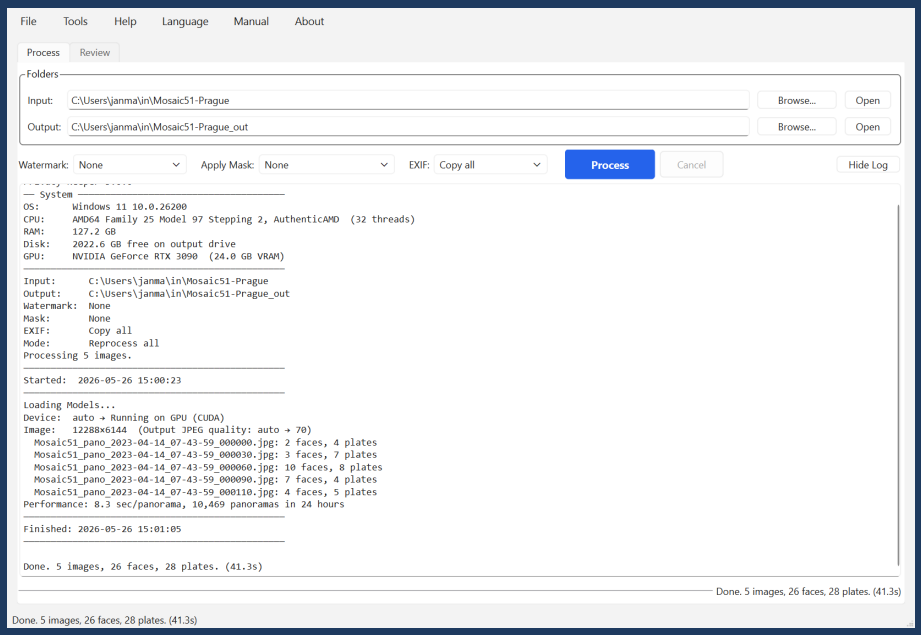

www.panohopper.com

Table of Contents

Table of Contents 2

Section 1 — Introduction 4

1.1 What Happens to Your Images 4

1.2 Supported Resolutions and Performance Data 4

Section 2 — System Requirements 5

Section 3 — Installation 6

3.1 Installing CUDA and cuDNN (for GPU acceleration) 6

3.2 Installing Privacy Keeper 6

Section 4 — Getting Started 7

4.1 First Launch 7

4.2 Selecting Folders 7

Section 5 — Settings 8

5.1 Face Detection Settings 8

5.2 Number Plate Detection Settings 8

5.3 Output Settings 9

5.4 Language 9

Section 6 — Processing Images 10

6.1 Starting a Job 10

6.2 Progress Bar and ETA 10

6.3 Cancelling 10

6.4 Completion 11

6.5 Log File 11

Section 7 — Reviewing Results 12

Section 8 — Edit Tab 14

8.1 Opening the Edit Tab 14

8.2 Editing Detections 14

Section 9 — Mask Editor 15

9.1 Drawing and Saving Masks 15

9.2 Applying Masks During Processing 15

Section 10 — Watermark 16

10.1 Tiled Watermark (Tools → Watermark) . . . 16
10.2 Applying the Watermark During Processing . . . 16
Section 11 — EXIF Editor . . . 17
11.1 Creating a Preset . . . 17
11.2 Applying a Preset During Processing . . . 17
Section 12 — GPU Acceleration . . . 18
12.1 Processing Speed . . . 18
Section 13 — Troubleshooting . . . 19
Section 14 — Support . . . 20

Section 1 — Introduction

Privacy Keeper automatically detects and blurs faces and license plates in equirectangular (360°) panoramic images. It is designed for photographers, mapping companies, and anyone who needs to anonymize panoramic imagery before publishing. Processing runs on your NVIDIA GPU for maximum speed, with automatic fallback to CPU.

1.1 What Happens to Your Images

Blurred images, log files and detections are all saved to the output folder. The log file (named `privacy_keeper_YYYY-MM-DD_HH-MM-SS.log`) records all processing details including system info, detection settings, per-image results, and performance metrics. The detections are saved to `detections.json`, a structured file you can use for further analysis or integration with other tools. Your original input images are never modified.

1.2 Supported Resolutions and Performance Data

Privacy Keeper is optimized for common panoramic cameras and resolutions:

Camera / Format	Resolution	GPU Time	CPU Time	FPS (GPU)
Mosaic 51	12,288 × 6,144 px	~2 sec	~26 sec	~0.5
Insta360 X4	8,192 × 4,096 px	~1.2 sec	~15 sec	~0.8
GoPro MAX	5,760 × 2,880 px	~0.7 sec	~9 sec	~1.4

These figures are measured on an NVIDIA RTX 3090 (GPU) and AMD Ryzen 9 7950X3D (CPU). Your performance may vary based on your hardware.

Section 2 — System Requirements

- Windows 10 or Windows 11 (64-bit)
- 8 GB RAM minimum, 16 GB recommended
- 4 GB free disk space
- For GPU acceleration: NVIDIA GPU with CUDA 12.x support (GTX 10xx or newer)
- For GPU acceleration: CUDA Toolkit 12.9.1 and cuDNN 9.20.0 installed

Section 3 — Installation

3.1 Installing CUDA and cuDNN (for GPU acceleration)

GPU acceleration is optional — the application falls back to CPU automatically if CUDA is not installed.

Step 1 — CUDA Toolkit 12.9.1

Download and install from: https://developer.nvidia.com/cuda-12-9-1-download-archive

Select: Windows → x86_64 → your Windows version → exe (local)

Step 2 — cuDNN 9.20.0

Download from: https://developer.nvidia.com/cudnn-9-20-0-download-archive

Select: Windows → x86_64 → Windows 11 → exe (local)

Run the installer and follow the on-screen instructions. No manual folder copying is needed.

3.2 Installing Privacy Keeper

Run `PrivacyKeeper_Setup_3.0.0.exe` and follow the on-screen instructions.

Section 4 — Getting Started

4.1 First Launch

Double-click the Privacy Keeper shortcut on your desktop or in the Start Menu. The main window opens with two tabs: **Process** and **Review**. Additional tools open as closeable tabs via the **Tools** menu: detection editing (**Tools** → **Edit**), tiled watermarks (**Tools** → **Watermark**), masks (**Tools** → **Mask**), EXIF presets (**Tools** → **EXIF**), and detection/output settings (**Tools** → **Settings**).

We strongly suggest processing known test data first to verify your installation. Download our test dataset:

www.panohopper.com/test-data/Fort_Myers_Beach_2022.zip

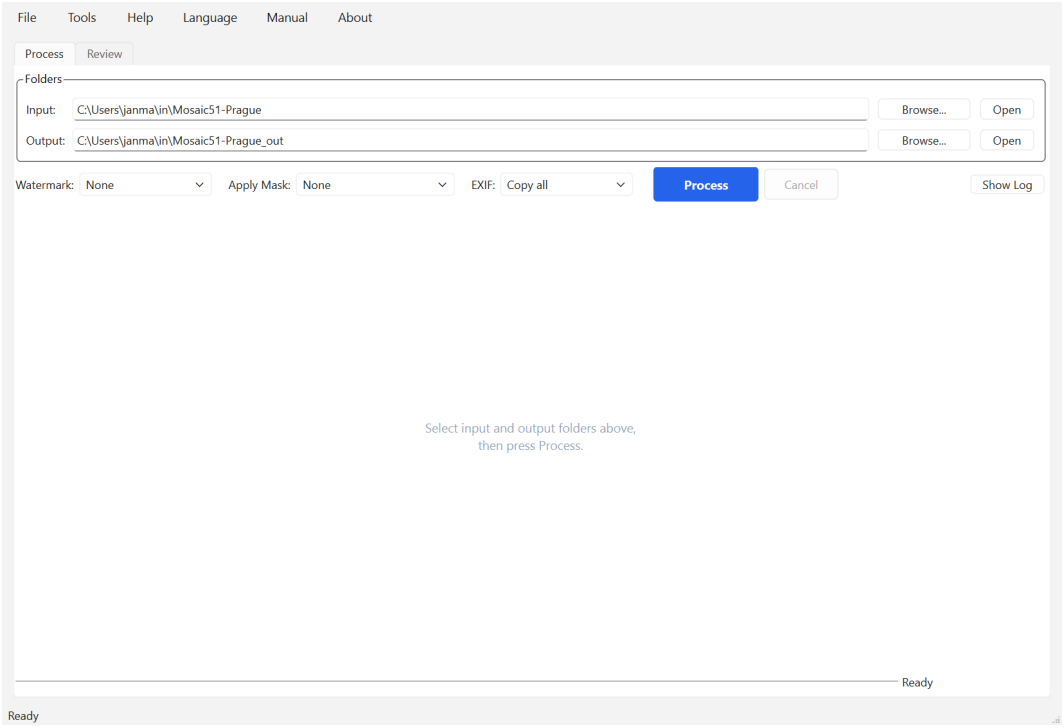

The main window on first launch. Select folders and press Process — no other steps needed.

4.2 Selecting Folders

Input Folder: Click **Browse…** next to the Input field and select the folder containing your panoramic JPEG images.

Output Folder: Automatically set to the input folder name plus `_out` whenever you pick a new input folder. You can override it at any time. The folder is created automatically if it does not exist.

The **Open** button next to each field opens the folder in Windows File Explorer.

Section 5 — Settings

Open **Tools** → **Settings** to adjust detection and output parameters. The Settings tab stays open until you close it.

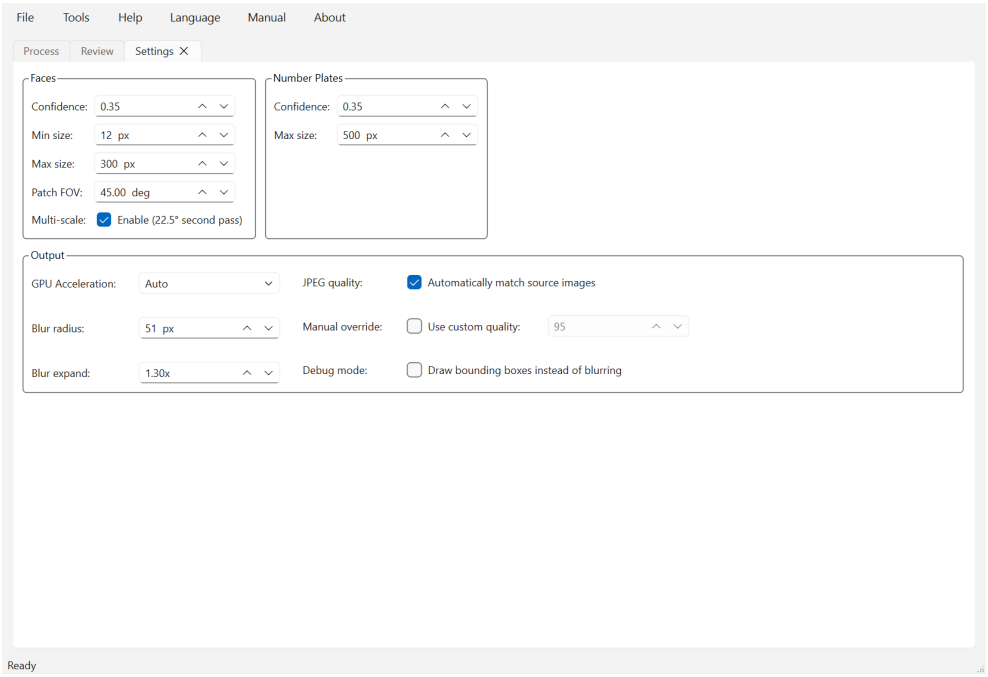

The Settings tab (Tools → Settings). Separate confidence thresholds for faces and plates.

5.1 Face Detection Settings

Setting	Default	Description
Face confidence	0.35	Minimum face detection score (0.1–0.99). Lower = more detections but more false positives.
Primary FOV	45°	Field of view for the rectilinear projection patches.
Multi-scale	Enabled	Runs a second detection pass at 22.5° FOV to catch very close faces.
Min face size	12 px	Faces smaller than this in the patch are ignored.
Max face size	300 px	Faces larger than this in the patch are ignored.

5.2 Number Plate Detection Settings

Setting	Default	Description
Plate confidence	0.35	Minimum plate detection score (0.1–0.99). Independent of the face threshold.
Max plate size	500 px	License plates with a bounding box larger than this are discarded as false positives.

5.3 Output Settings

Setting	Default	Description
Blur radius	51 px	Gaussian blur kernel size. Larger = stronger blur. Must be odd.
Blur expand	1.3x	Expands the detected bounding box before blurring.
JPEG quality	Auto-match	Detects and matches each source image's original JPEG quality. Override with a fixed value (50–100) if needed.
Debug mode	Off	Draws coloured bounding boxes instead of blurring (red = face, green = plate).
GPU Acceleration	Auto	Auto: uses GPU if CUDA is available, otherwise CPU.

5.4 Language

Click **Language** in the top menu bar to switch between English, Deutsch, and Français. The choice is saved on next launch.

Section 6 — Processing Images

6.1 Starting a Job

1. Select your input and output folders.
2. Adjust settings if needed (defaults work well for most cases).
3. Click **Process**.

If the output folder already contains images with the same names, a dialog appears:

- **Continue (skip existing)** — Only processes new images.
- **Reprocess All** — Overwrites all output images.

The log starts with your PC's system specifications, then confirms GPU or CPU is active. Each image is logged with face and plate counts.

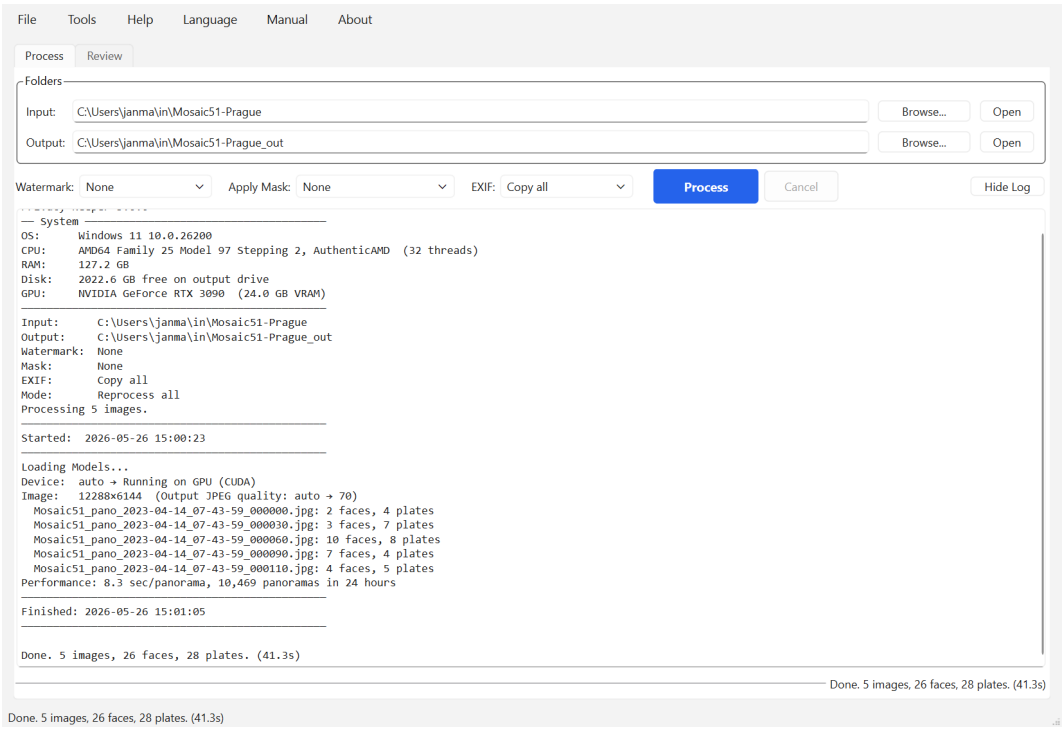

The Process tab after a completed job, showing the full log with system info, per-image results, and performance summary.

6.2 Progress Bar and ETA

The progress bar shows current / total images. An estimated finish time appears once the first image completes, using a rolling 10-image average.

6.3 Cancelling

Click **Cancel** at any time. The current image finishes before the job stops. Already-processed images remain in the output folder.

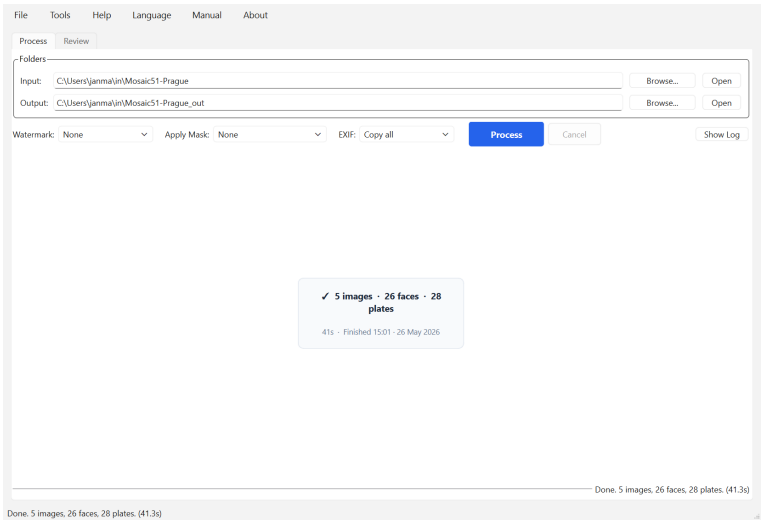

The summary card shown when the log is hidden after a completed job.

6.4 Completion

When processing finishes the **Review** tab opens automatically. A summary card on the **Process** tab shows totals and elapsed time.

6.5 Log File

A timestamped log file is saved to the output folder after each job. The **Open Log File** button in the Process tab toolbar toggles the log pane.

Section 7 — Reviewing Results

The **Review** tab shows processed images one at a time at full resolution. Use the scroll wheel to zoom and drag to pan.

Control	Action
< Prev (P) / Left arrow	Go to previous image
Next (N) > / Right arrow	Go to next image
Fit	Reset zoom to fit the window
Show Detections	Overlay bounding boxes from detections.json (red = face, green = plate).
Scroll wheel	Zoom in / out
Click and drag	Pan the image

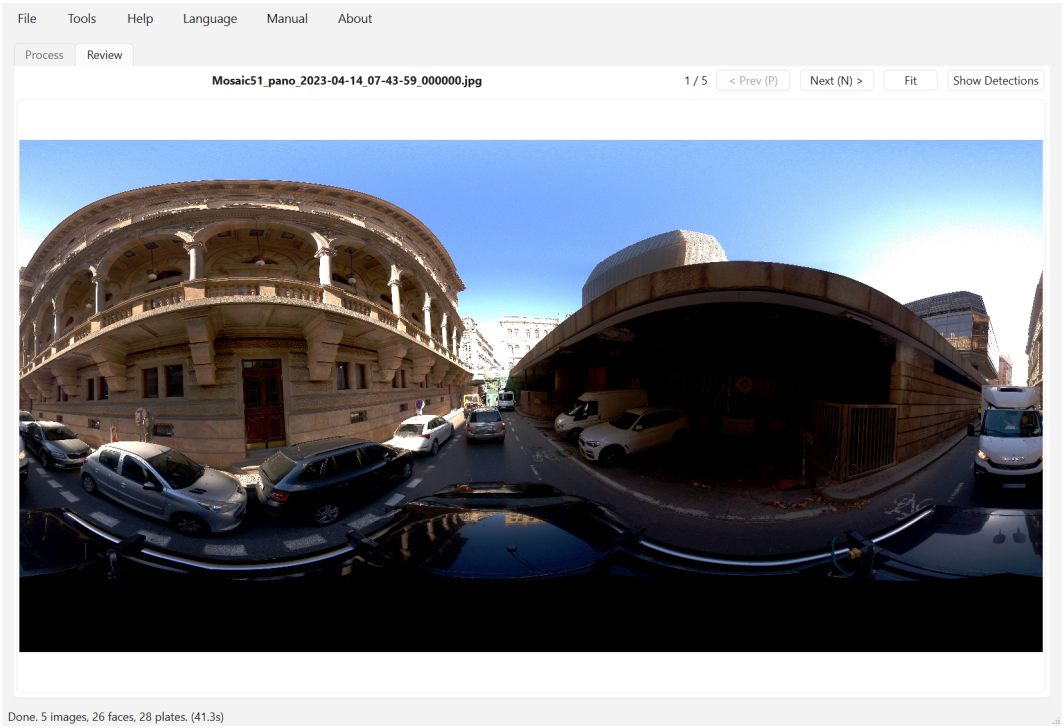

The Review tab showing a blurred panorama at full resolution.

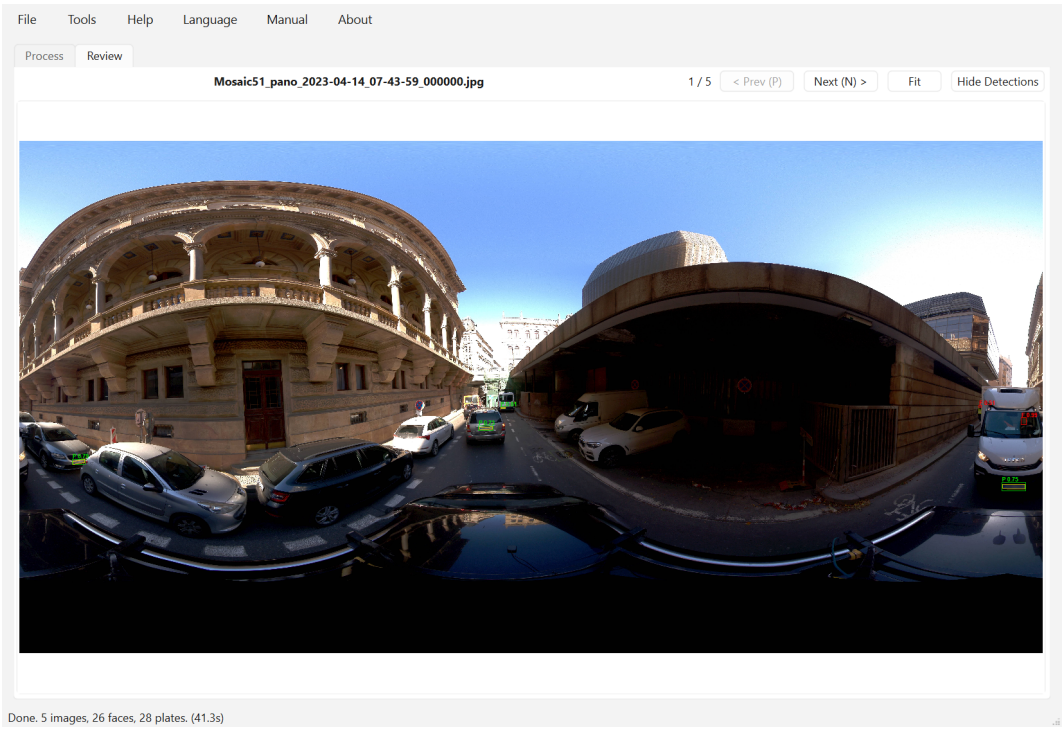

Detection overlay: red boxes = faces, green boxes = license plates.

The Review tab opens automatically when processing finishes and reloads the output folder on every switch during processing.

Section 8 — Edit Tab

The **Edit** tab (opened via **Tools** → **Edit**) lets you manually correct detection results after a batch has run. Add missing detections, remove false positives, and reprocess individual images.

8.1 Opening the Edit Tab

Go to **Tools** → **Edit**. Images and detections load automatically from `detections.json` if an output folder is set.

8.2 Editing Detections

Navigate with **‹ Prev** / **Next ›** or arrow keys. Existing boxes are shown as coloured overlays.

- **Add Face** — drag a rectangle to add a face detection.
- **Add Plate** — same but marks a license plate.
- **Right-click a box** — deletes that detection.
- **Undo** — reverts the last add or delete.
- **Revert** — resets all edits for the current image.
- **Save** — writes edits back to `detections.json`.
- **Reprocess** — re-blurs the current image using the edited detections.

Section 9 — Mask Editor

The **Mask** tab (opened via **Tools** → **Mask**) allows you to draw custom polygon masks on panoramic images. Masked regions are blurred during processing in addition to detected faces and plates.

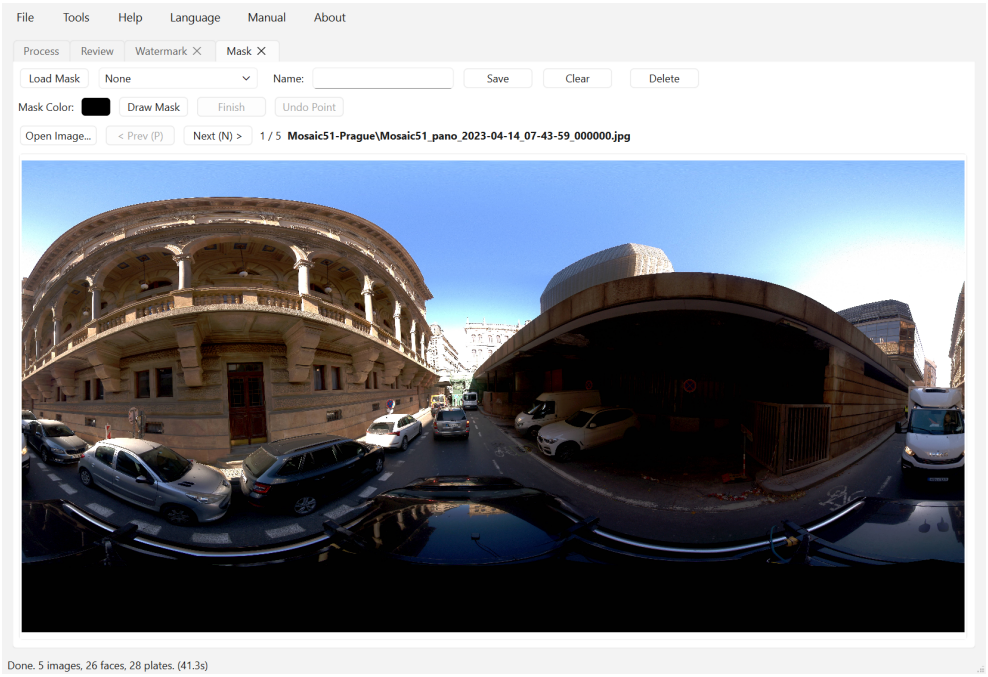

The Mask Editor tab (Tools → Mask). Draw polygon masks and save them as named presets.

9.1 Drawing and Saving Masks

1. Select a mask colour (default: black).
2. Click on the image to place vertices and draw a polygon.
3. Enter a name in the Mask Name field.
4. Click **Save Mask**.

9.2 Applying Masks During Processing

Select a saved mask from the **Apply Mask** dropdown in the Process tab toolbar before clicking Process. Masks are saved in `~\.panohopper\masks\`.

Section 10 — Watermark

The tiled watermark tool composites repeating text onto output images after blurring. Open it from **Tools → Watermark**.

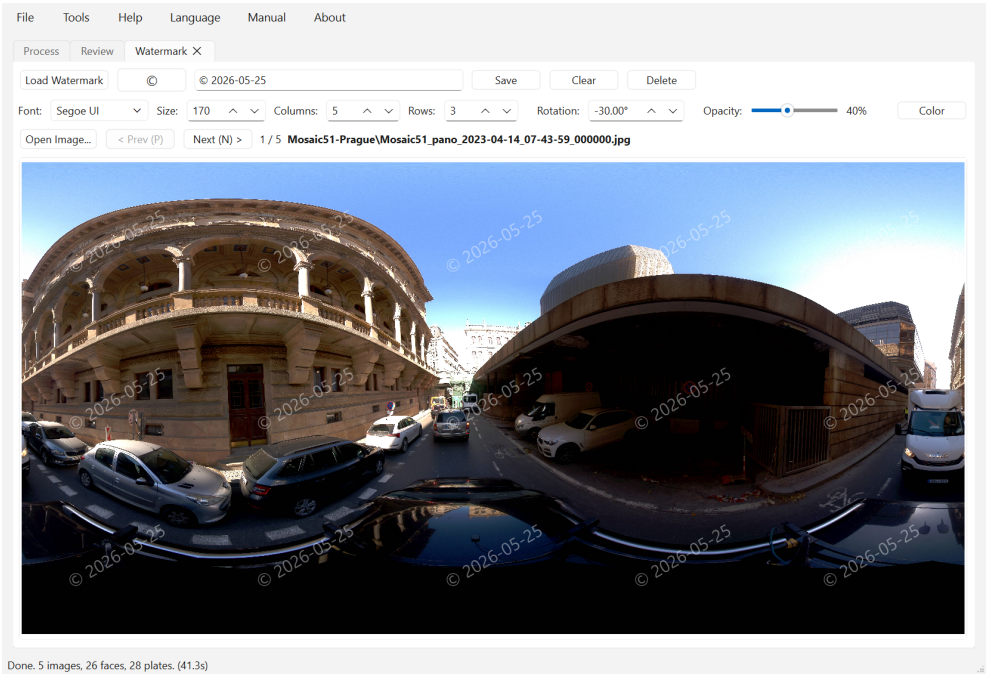

The Watermark tab (Tools → Watermark). Configure text, font, grid density, rotation, and opacity.

10.1 Tiled Watermark (Tools → Watermark)

Repeats text in a regular grid across the entire image. Useful for ownership notices or confidential stamps.

1. Enter the watermark text.
2. Choose a font and set size and opacity.
3. Adjust tile spacing and rotation angle.
4. Enter a name and click **Save**.

10.2 Applying the Watermark During Processing

Select the desired preset from the **Apply Watermark** dropdown in the Process tab toolbar before clicking Process. Presets are saved in `~\.panohopper\watermarks\`.

Section 11 — EXIF Editor

The **EXIF** tab (opened via **Tools** → **EXIF**) lets you create presets that control which metadata tags are copied, stripped, or edited on output images. By default, Privacy Keeper copies all EXIF metadata from each source image — GPS coordinates, camera model, capture time, and all other tags are preserved unchanged.

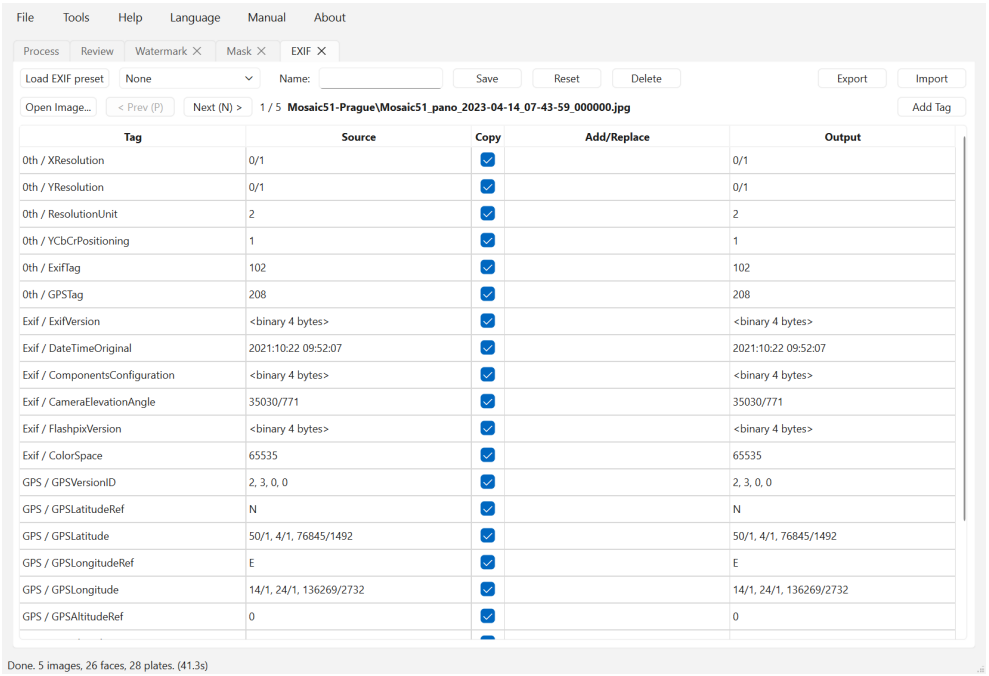

The EXIF Editor tab (Tools → EXIF). Browse metadata tags and build strip/edit presets.

11.1 Creating a Preset

In the EXIF tab:

1. Images from the input folder load automatically when the tab opens.
2. Browse the EXIF tags read from the current image.
3. Select tags to **strip** (remove entirely) or enter new values to **edit** them.
4. Enter a name and click **Save** to store the preset.

11.2 Applying a Preset During Processing

The toolbar at the top of the **Process** tab includes an **EXIF** dropdown. The default is **Copy all** — all source metadata is preserved. Select a saved preset from the dropdown before clicking **Process** to apply your tag edits or strips to every output image in the batch.

Presets are saved as JSON files in `~\.panohopper\exif\` and persist across sessions.

Section 12 — GPU Acceleration

Privacy Keeper uses your NVIDIA graphics card to speed up detection significantly. The **GPU Acceleration** setting in the Settings tab controls this:

- **Auto** — uses GPU if CUDA is installed and available, otherwise CPU (recommended).
- **GPU** — forces GPU; shows an error if CUDA is not available.
- **CPU** — forces CPU regardless of GPU availability.

If an NVIDIA GPU is detected but CUDA or cuDNN is missing, the Process log displays a warning with installation instructions.

12.1 Processing Speed

Measured on an NVIDIA RTX 3090 (GPU) and AMD Ryzen 9 7950X3D (CPU), processing 12 288 × 6 144 pixel Mosaic 51 panoramas:

Mode	Hardware	Time per image	Panoramas per day
GPU	NVIDIA RTX 3090 (24 GB VRAM)	~2 seconds	~43,000
CPU	AMD Ryzen 9 7950X3D (16 cores)	~26 seconds	~3,320

GPU processing is approximately 13× faster than CPU on this hardware.

Section 13 — Troubleshooting

No images found

Make sure the input folder contains `.jpg` or `.jpeg` files directly (not in subfolders).

GPU not used / CUDA warning

Open **Help** → **CUDA Installation Guide** and install CUDA 12.9.1 and cuDNN 9.20.0. Use CPU mode in the meantime.

Faces or plates not detected

Lower the Confidence threshold (e.g. to 0.20). Enable Multi-scale. Ensure images are equirectangular panoramas.

Too many false detections

Raise the Confidence threshold (e.g. to 0.55). Use Debug mode to inspect detections.

Output file size differs

Re-enable JPEG auto-match in Settings. All EXIF metadata (GPS, camera, date/time) is always preserved in output.

Section 14 — Support

Website: https://www.panohopper.com