Blood Stain Classifier - Getting Started Guide

Version: 2026-06-03

Administrator contact: daniel@attinger-consulting.com

What does this app do?

You upload a photograph of a transfer stain. The app tells you:

- **What object** made the stain (one of 21+ trained objects, e.g. Knife Blade, Bare Finger, Cotton Sock), shown as the top few most-likely candidates.
- **Which direction** the object moved, drawn as an arrow on your image and reported as an angle.
- **A heatmap** showing which pixels the AI used to decide.

Each result comes with a confidence score from 0% to 100%. Treat anything below 98% with caution.

Terminology used throughout this guide follows **ASB Technical Report 033, 1st Ed. 2017** (Terms and Definitions in Bloodstain Pattern Analysis). See the Glossary at the end for full definitions.

> **New since the May workshop.** Several things changed after the workshop, and they are reflected throughout this guide: photos are now automatically turned to landscape so phone pictures taken over a table no longer give random results; the app now warns you before spending credits if the stain is framed too small; registration now asks for your name and organisation; the object result is now shown as the top three candidates; submission modes now have clear data-retention rules; and the workshop quiz now shows a scale bar on each stain. Online credit purchase is not yet live - see Step 2.

Under the hood

The app combines several stages of image processing and machine learning. The pipeline is the result of multi-year proprietary development and tuning; the high-level outline below is provided so you can interpret the results, not so you could rebuild it.

1. Orientation and segmentation. When your image is uploaded it is first turned to landscape (long edge horizontal) and resized to a standard manageable size, then cleaned: lighting, colour balance and contrast are normalised so that photographs from different cameras and lighting conditions become comparable. The stain itself is then isolated from the background using a custom multi-method segmentation procedure, so that subsequent stages look only at blood pixels.

2. Feature extraction. A deep convolutional neural network, fine-tuned on our own data, converts the segmented stain into a numerical fingerprint that captures shape, texture and edge geometry. Two complementary regions of the stain are fingerprinted independently to preserve directional information.

3. Rotation analysis. For the Detailed Analysis level, the stain is examined at several orientations. Each orientation gives a partial directional vote, weighted by the model's confidence for that orientation. These votes are combined using a proprietary vector-integration scheme that yields the final swipe angle at a finer resolution than any single tested orientation.

4. Classification. A two-output classifier is run on the fingerprint: one output describes the object that made the stain (selected from the 21+ trained classes), the other identifies the direction of relative motion. What the classifier learns about an object also informs its directional judgement.

5. Confidence estimation. Each prediction is accompanied by a confidence score derived from the classifier's output distribution. For the direction estimate in Detailed mode, an additional consensus score reflects how well the votes from different orientations agreed. Low consensus is a deliberate signal of uncertainty rather than a number you should round up.

Training data. The model was trained on the order of thousands of carefully curated photographs collected and annotated in-house, covering every supported object across both swipe directions, plus comparable numbers of wipes for objects where wipes are realistic. Human blood was used to generate the training data. Each image was manually verified for ground truth. Curation effort, not algorithm choice alone, is what makes the system reliable.

Before you start

You only need a web browser and the URL your administrator gave you. The app runs entirely in the cloud - nothing to install.

Photographing the stain well matters more than anything else. A good photo is on a plain white background (cardstock, paper, or wall), tightly cropped around the stain, well-lit and in focus.

Two things the app now does for you, both added after the workshop:

- **Orientation is fixed automatically.** At the workshop, people who photographed stains lying flat on a table (phone held nearly flat, pointing down) often got pictures that came out taller than wide, which produced random answers. The app now always turns your picture to landscape (long edge horizontal) before analysing it, so you no longer have to worry about how you held the phone.
- **The app checks your framing.** If the stain fills less than about 40% of the frame, the app pauses before spending any credits and suggests you retake the photo closer, aiming for the stain to fill roughly 80% of the frame. This is just a gentle warning for very loosely framed photos - you can still choose to proceed, but tighter framing gives more reliable results.

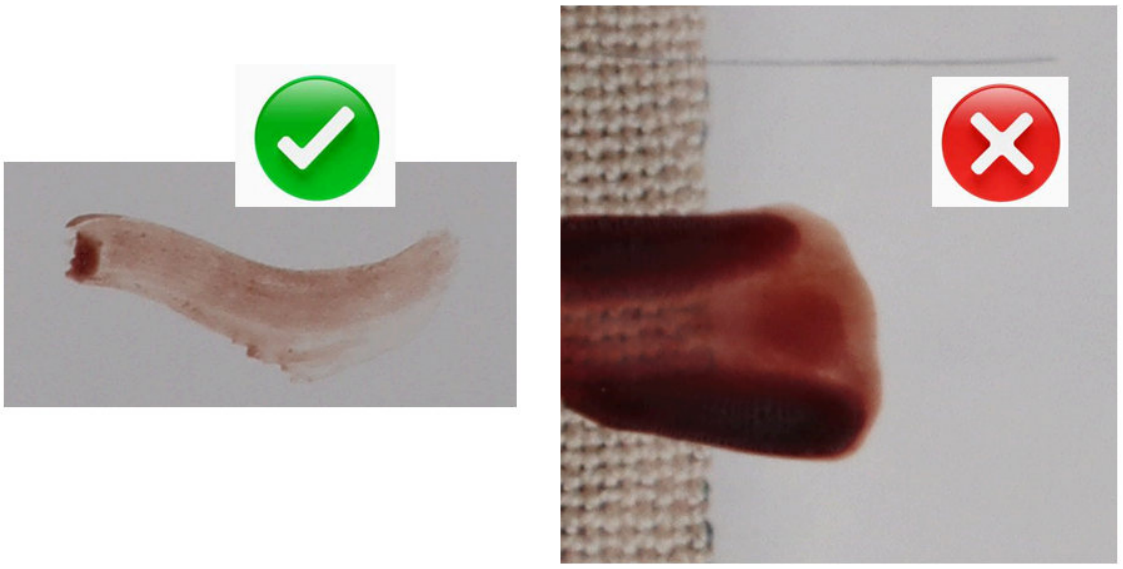

FIGURE 1 - Good vs bad photo. Side-by-side: left = a clean transfer stain on white background, well-cropped; right = the same or similar stain on a cluttered/textured background, awkward angle, or blurry.

Step 1 - Open the app and log in

Open your browser and go to the URL your administrator sent you. The login screen appears.

If you already have an account, enter your email and password.

If you are new, click **Register** and fill in:

1. **Email**
2. **Full Name** (so the administrator knows who the account belongs to)
3. **Organization** (e.g. your laboratory or police department)
4. **Password** (at least 8 characters with uppercase, lowercase, and numbers)

Creating an account means you agree to the **Data Retention Policy** (linked under the form - see "Data and image retention" below).

Your account starts in **Pending** state - an administrator must approve it before you can classify images. Approval also grants your initial welcome credits. If approval takes more than a working day, email daniel@attinger-consulting.com.

If you have a **voucher code**, enter it during registration, or later via the **Redeem Voucher** button in the top bar. Vouchers add credits.

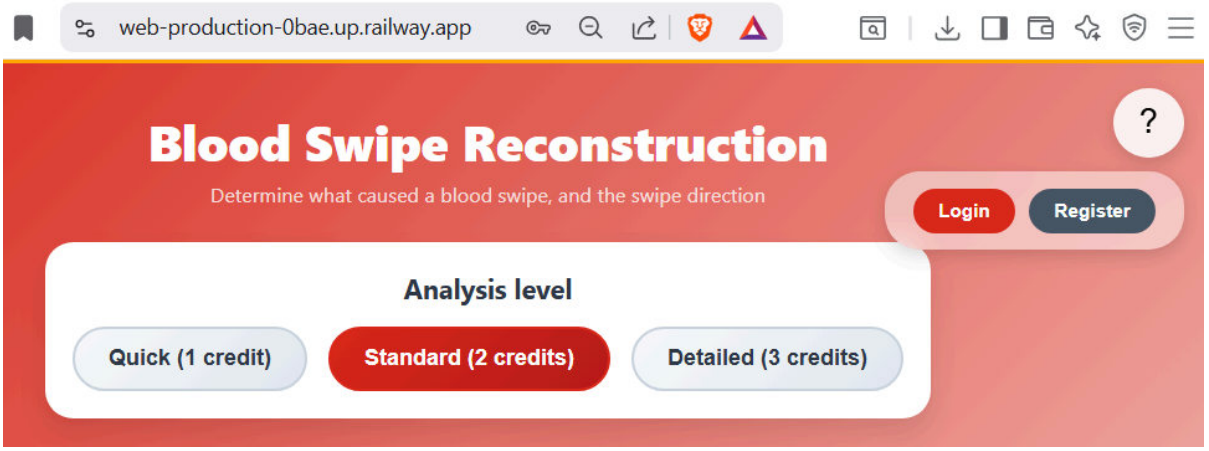

FIGURE 2 - Screenshot of the login / register modal as it appears in the browser, showing the Email, Full Name, Organization, and Password fields.

Step 2 - Credits

Each classification consumes credits. Your balance is shown in the top bar after login.

Online credit purchase is not yet available. The **Credits** badge opens a "Credit Purchase - Coming Soon" notice. For now, you get credits in one of two ways:

- **Redeem a voucher code** (during registration, or via **Redeem Voucher** in the top bar). Credits are added immediately.
- **Request credits by email** at daniel@attinger-consulting.com if you need more.

How many credits a classification costs depends on the analysis level and submission mode you pick - see Step 3.

Step 3 - Pick an analysis level and a submission mode

Two rows of buttons appear under the header.

Analysis level controls how thoroughly the image is analysed:

Button	Base cost	What it does
Quick	1 credit	Fastest. Single classification pass.
Standard *(default)*	2 credits	Adds a heatmap. Recommended.
Detailed	3 credits	Adds rotation analysis to find the best swipe angle.

Submission mode controls whether your image is stored, who can see it, and how long it is kept. The base cost above is then adjusted by the mode:

Button	Cost	What happens to your image
Private	2x base (2 / 4 / 6)	Image is **never stored** - it is processed in memory only. Neither the administrator nor anyone else can see your image afterward. Use for

Button	Cost	What happens to your image
		sensitive or confidential casework.
Standard *(default)*	Base (1 / 2 / 3)	Image is stored so you can review it, then **automatically erased 30 days after submission**. Result and metadata stay in your history. Result visible to the administrator.
Research	Half base, minimum 1 (1 / 1 / 2)	You contribute the image plus ground-truth details (entered after classification). Kept **3 years** (longer if used to train the model).

A few things to know about **Research mode**, which is new in its current form:

- It is a **data-contribution programme**, not just a discount. The reduced price is in exchange for accurate, good-faith ground-truth information you provide.
- After the image is classified, a **Research Ground Truth** form asks you for the real object type, swipe class, blood source, swipe direction, and surface. Enter genuine, best-effort answers only.
- In research mode you submit **one image at a time** (the ground-truth step is per image).

The selected button in each row is highlighted. A short line under the buttons reminds you what the current mode does.

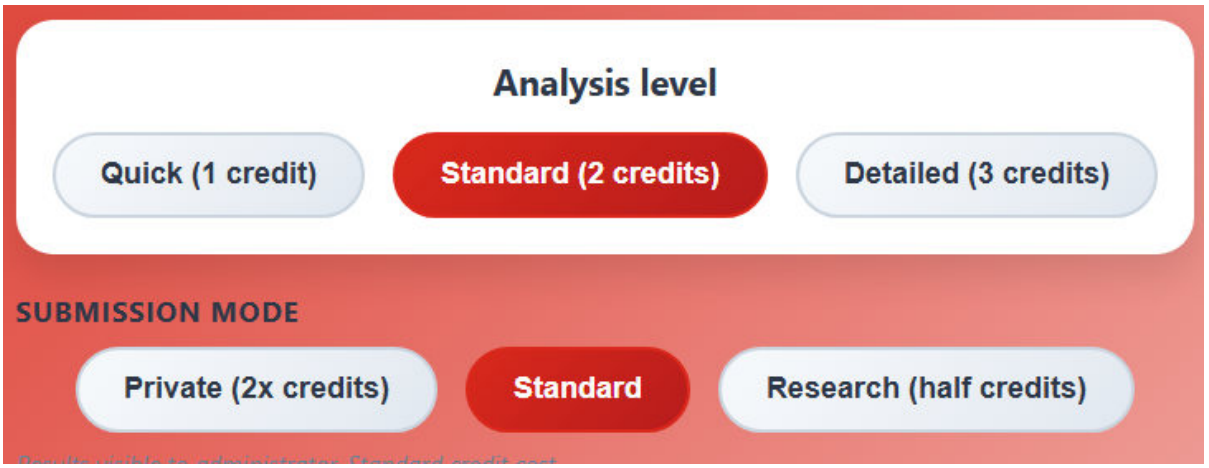

FIGURE 3 - Cropped screenshot of the two button rows: "Analysis level" (Quick / Standard / Detailed) with Standard highlighted, and "Submission Mode" (Private / Standard / Research) with Standard highlighted, plus the description line beneath them.

Step 4 - Upload your image

Drag your file into the **drop zone** in the centre of the page, or click **Browse/Take Photo**. On a phone, this lets you take a photo directly with the camera.

Accepted formats: PNG, JPG, JPEG. In Standard and Private modes you can upload several files at once - they are processed one after another, each with its own progress bar. In Research mode you upload one image at a time.

Two reminders, both from the workshop fixes:

- You do not need to rotate your photo - the app turns every image to landscape automatically.
- If the stain is framed too small, the app will pause and suggest a closer retake before any credits are spent.

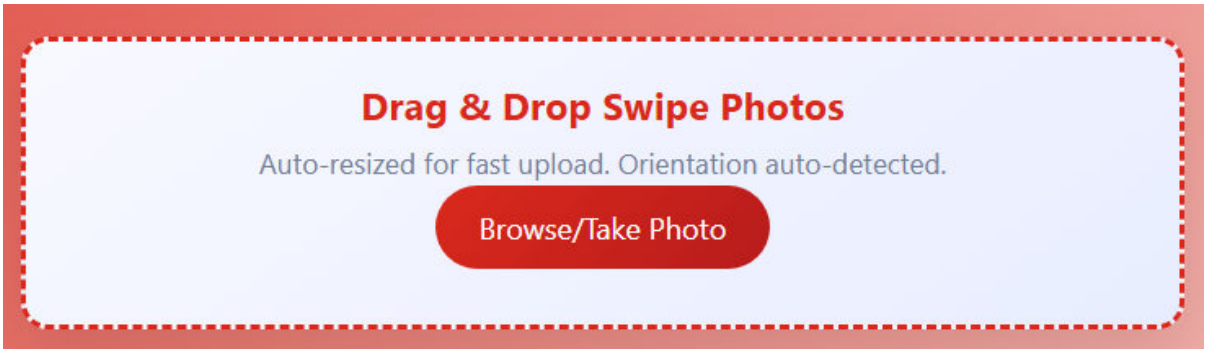

FIGURE 4 - Screenshot of the drop zone mid-upload, showing one image with a partial progress bar.

Step 5 - Read the results

When processing finishes, a result card appears for each image. The card shows:

1. **Your photo**, with a direction arrow drawn on top, reported as an angle with a tolerance (for example "39deg +/- 45deg"). Near-square photos may show white bars on the sides - this is correct, it is exactly the landscape frame the classifier analysed.
2. **Bloody Object** - now shown as the **top three candidates** rather than a single guess. A small pie shows the share of each candidate, and each candidate's reference photo is displayed beneath it with the object's name. The most likely object is highlighted. The pie slice sizes convey the margin between candidates; the headline candidate still carries a confidence percentage.
3. **Swipe Direction** - the angle of swipe direction, with a confidence percentage.
4. **Thresholded Image** - the AI attempt to extract the bloody region (highlighted in red) from the background. Background is ignored for classification.
5. **Heatmaps** - a colour overlay showing where the AI looked. Red and orange areas weighed most; blue areas were ignored.

If the heatmap lights up the background rather than the stain, the image is probably not reliable.

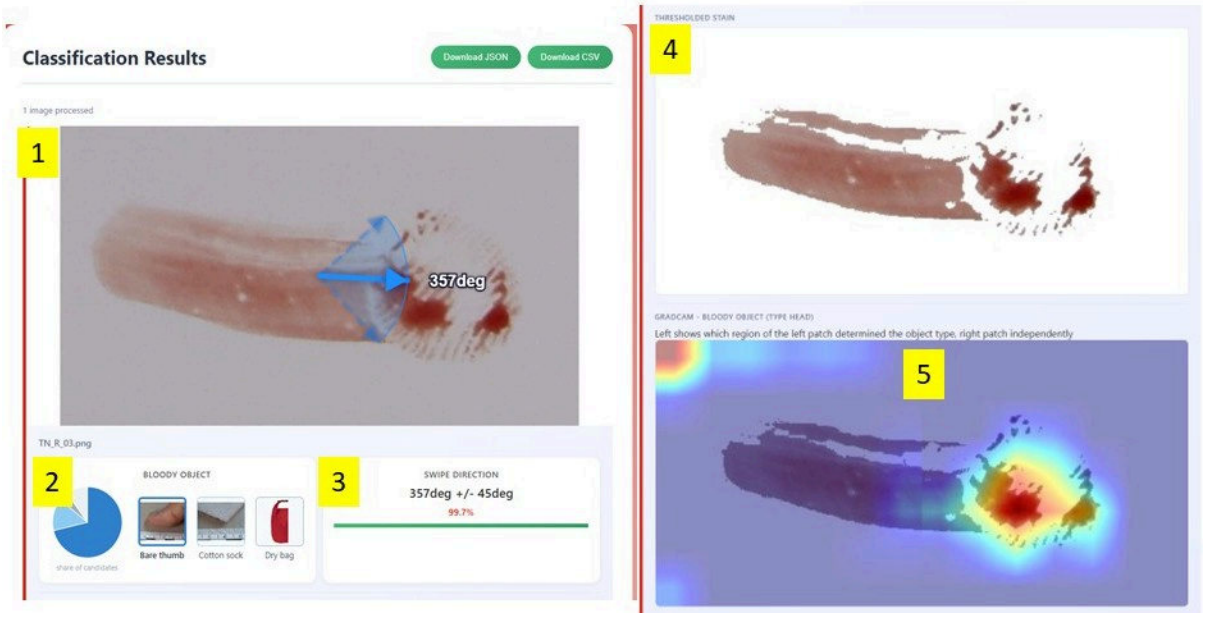

FIGURE 5 - Annotated result card, showing the direction arrow, the top-three object candidate pie with reference photos, the thresholded image, and the heatmaps.

Step 6 - Export

Two buttons under the results download everything:

- **Download JSON** - full structured data, all confidence scores.
- **Download CSV** - one row per image, Excel-friendly.

Remember that in Standard mode the **images** are erased 30 days after submission, so download anything you need to keep within that window.

Past classifications

Click **History** in the top bar to browse or download every image you have submitted, with predictions and timestamps.

Note on what stays available:

- **Standard mode:** the result and metadata stay in your history permanently, but the stored images are removed 30 days after submission.
- **Private mode:** the result and metadata are recorded, but there were never any images to view.
- **Research mode:** images and your ground-truth entries are kept for 3 years (longer if used to train the model).

Data and image retention

Because the app can store forensic images, there is a written **Data Retention Policy** (effective 2 June 2026). You agree to it when you create an account, and a link to it appears under the registration form. The essentials:

- **Standard mode:** images kept 30 days, then automatically erased; results stay in history.
- **Private mode:** images never stored.
- **Research mode:** images and ground truth kept 3 years, and potentially indefinitely where a submission becomes part of a training or validation dataset (in anonymised or aggregated form).

You can request early deletion of a specific submission's images, request a copy of your data, correct your account details, or close your account (which removes your account data and associated records). For any of these, or any questions, contact daniel@attinger-consulting.com.

Forensic images may constitute personal or sensitive data in some contexts. You are responsible for having a proper basis to upload an image, and for choosing **Private mode** when confidentiality is required.

What the app recognises

The model was trained on transfer stains made by **21 different objects (and counting)**. The exact list can be retrained over time - click **History** after submitting a few images to see the current set. The objects fall into broad groups: tools (knife blade, scissor blade, screwdriver, drill bit, key), bodily (bare finger, bare thumb, hair, latex glove), fabric (cotton sock, sports shorts, leather band, two fabric types), paper (paper towel, toilet paper, crumbled paper), and other household items (sponge, Q-tip, plastic bag, microscope slide, paint brushes).

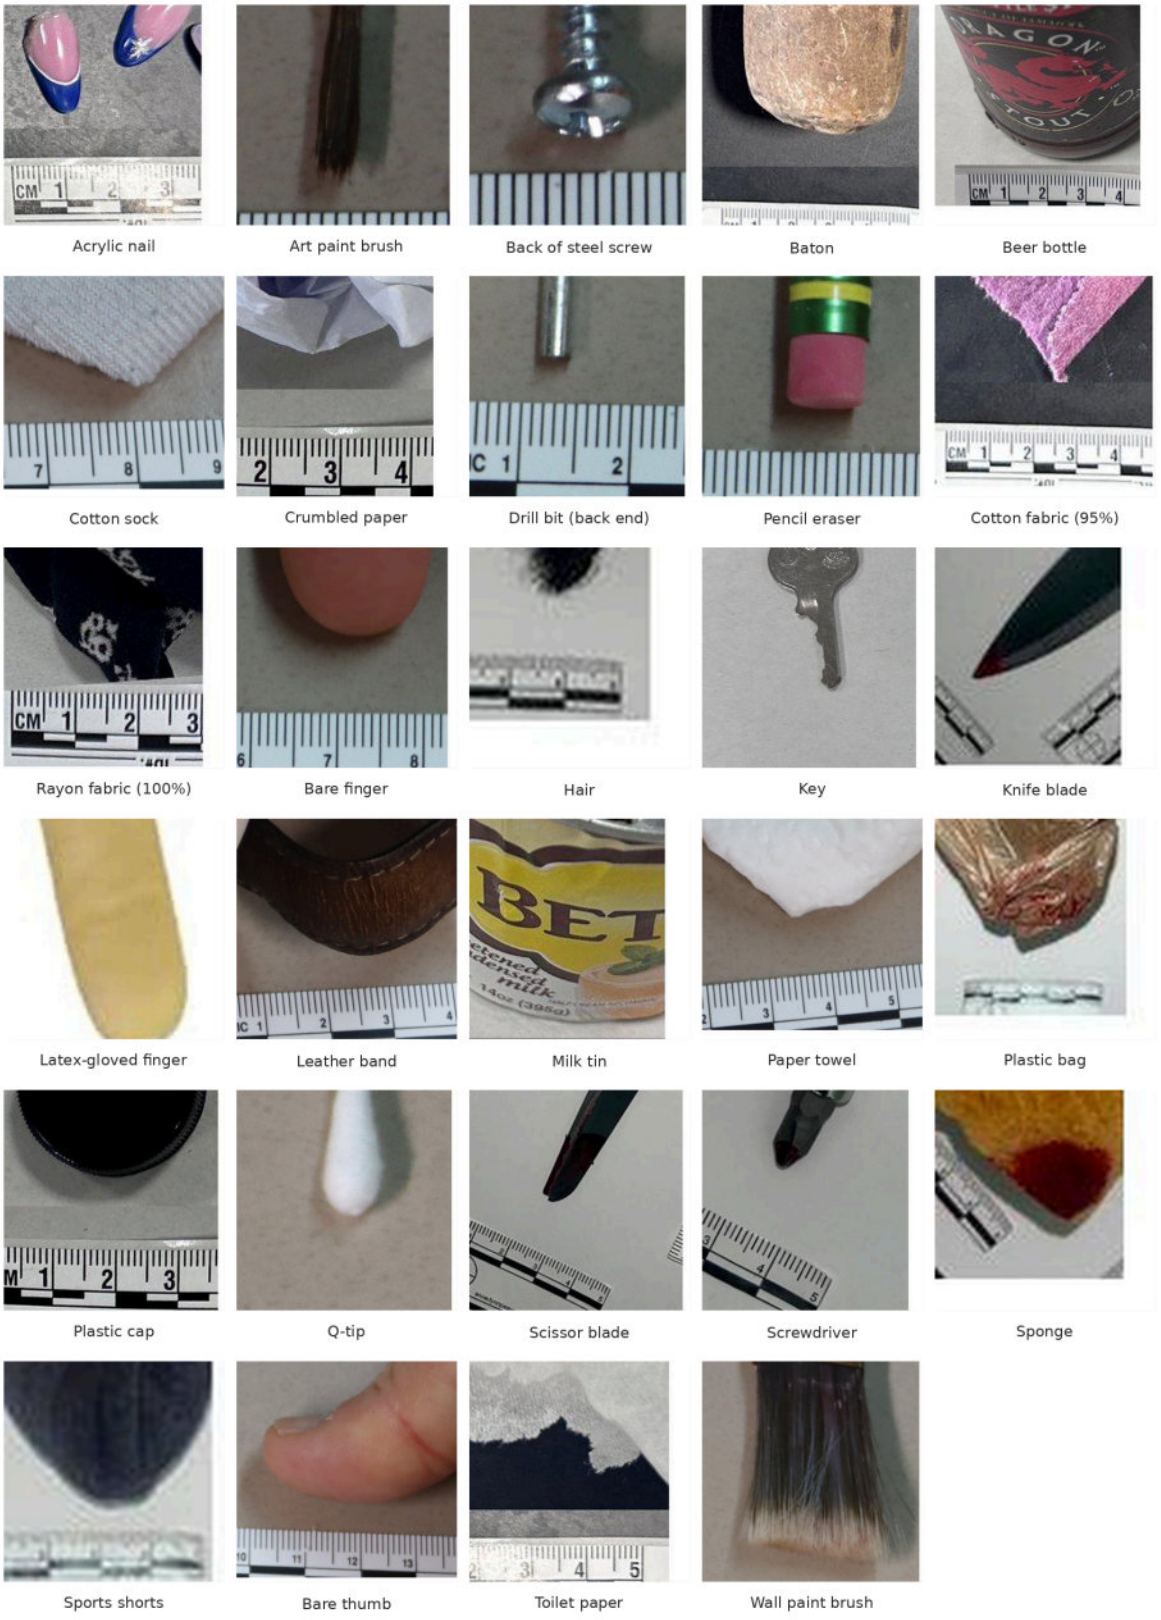

FIGURE 6 - A grid of the object reference photos with object names underneath each.

Quick troubleshooting

Symptom	What to do
Page does not load	Double-check the URL. If it was working earlier, the service may be restarting; wait a minute and refresh. If it still does not load, contact support daniel@attinger-consulting.com.
"Classifier not loaded" error	A model file is missing on the server. Contact support daniel@attinger-consulting.com.
Account stays in "Pending"	Waiting for admin approval. Contact support if approval does not occur as fast as you wish daniel@attinger-consulting.com.
"Retake suggested" before upload	The stain fills less than about 40% of the frame. Retake closer so it fills roughly 80%, then upload again. You may also proceed if you prefer.
Photo looked portrait but came out landscape	This is intended - the app always turns images to landscape so phone photos taken over a table classify correctly.
White bars on the sides of the result image	Normal for near-square photos; it is the landscape frame the classifier sees, not an error.
Confidence below 98%	Retake the picture, tighter crop, better light. If it does not improve, use with caution or contact support.
Out of credits	Redeem a voucher code, or email support for more. Online purchase is not yet available.
Anything else	Email support at daniel@attinger-consulting.com.

Workshop quiz (optional)

If the administrator has enabled Workshop Mode, the app also runs an anonymous quiz at `<your-app-url>/workshop/`. Participants enter a nickname (no login needed) and see a series of transfer-stain images. For each image they answer:

1. **Wipe** or **Swipe**?
2. **To Left** or **To Right**?
3. **Which object** made it? (3 options shown)

Each stain image now carries an **adaptive scale bar** (for example a "1 cm" marker) so participants can judge the real size of the stain - added in response to workshop feedback that the quiz stains had no scale. The bar appears only when scale information is available for that image.

After each answer the app reveals the ground truth and the AI's prediction side-by-side. A final scoreboard compares the participant against the model.

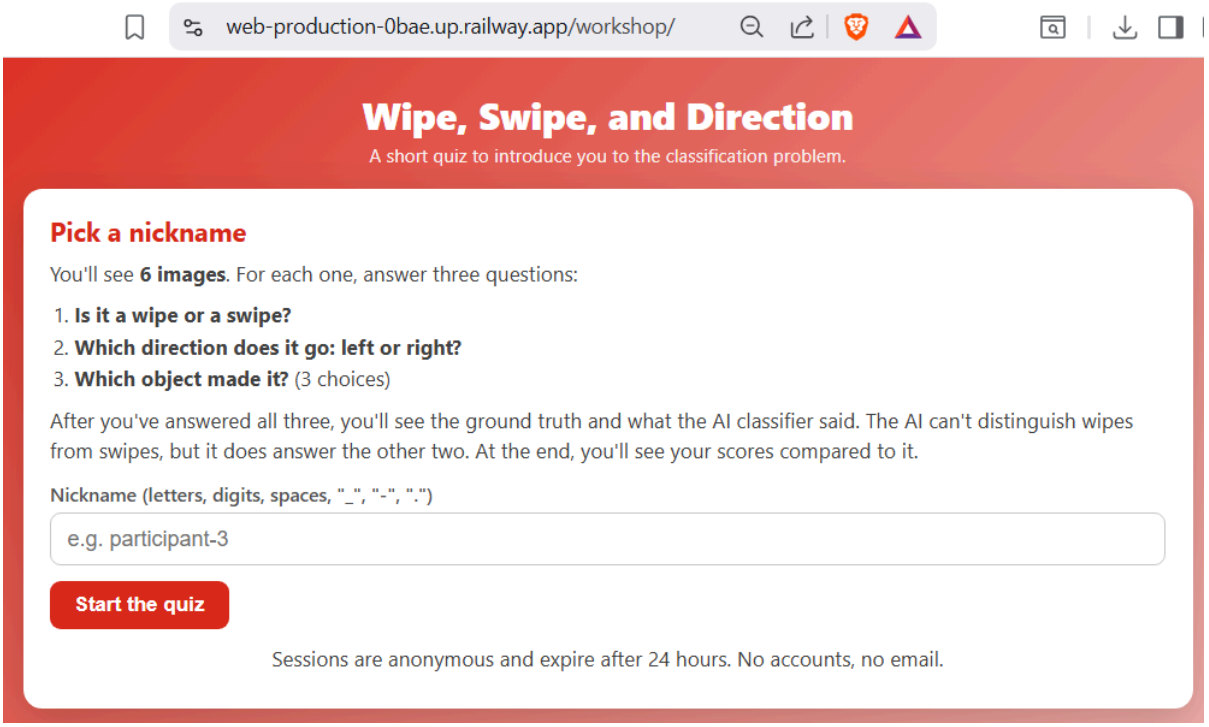

> *[FIGURE 7 - Screenshot of the workshop quiz screen: transfer-stain image with a scale bar on the left, the three question groups (Wipe/Swipe, Direction, Object) on the right.]*

Glossary

Definitions below follow **ASB Technical Report 033, 1st Ed. 2017** (Terms and Definitions in Bloodstain Pattern Analysis), AAFS Standards Board.

Transfer stain - A bloodstain resulting from contact between a blood-bearing surface and another surface.

Swipe - A bloodstain resulting from the transfer of blood from a blood-bearing surface onto another surface, with characteristics that indicate relative motion between the two surfaces.

Wipe - An altered stain resulting from an object moving through a preexisting wet bloodstain.

Confidence - how certain the model is, expressed as a percentage.

Grad-CAM heatmap - a visualisation showing which pixels of your image the neural network weighted most heavily.

Top-three candidates - the three most likely objects for a stain, shown together as a pie with reference photos, rather than a single best guess.

Submission mode - your choice of Private, Standard, or Research, which sets the credit cost and how long (if at all) your image is stored.

Scale bar - a small overlaid marker on a quiz stain (for example "1 cm") indicating the real-world size of the image, derived from the recorded dots-per-cm for that image.

Credits - units consumed each time you submit an image.

Voucher - a code that adds credits to your account, typically issued by an administrator.