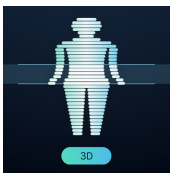

aPrehliadky — user guide

3D documentation of bodies, persons, objects and scenes with LiDAR · version 3.0 · forensika.eu

1. What you get and what the app needs from you

aPrehliadky builds a 3D model of a deceased body, a living person, an object or a scene from ordinary iPhone frames. With LiDAR every frame gets depth and gravity, the model is built directly in metres and can be measured. Everything runs on the phone: no cloud, no account, no upload, no AI-completed geometry.

Output: a virtual examination, findings anchored to the surface, measurements (distance, path along the surface, area, volume), surroundings from LiDAR (room mesh), a PDF report, a 3D documentation certificate (class A–D) and an .f3d package with hashes as an attachment for the case file.

Three things decide the quality: 1) the calibration target in the frame, 2) a slow, complete walk around the subject at three heights, 3) even light without flash. Without the target a measurement cannot be independently defended.

The app is a documentation tool. It supports professional judgement and does not replace it. It is not a certified measuring instrument — accuracy depends on capture conditions and is verified with a reference. The user is responsible for the lawfulness of capture and handling of the documentation.

2. Before capture — checklist

1. Device: iPhone Pro or iPad Pro with LiDAR, iOS 18+, charged to at least 50 % (processing takes minutes and heats the phone).
2. Calibration target T3 printed at 100 % on A4 landscape (scan list → menu … → Calibration target T3 for printing), mounted on a rigid board. Measure the circle centres A–B (247.00 mm) with a caliper and enter the value on the scan screen (Target print check). The app finds the circular fiducials in every frame even from 2 m, triangulates them directly in the model after processing and compares 10 distances with the reference — a model scale independent of LiDAR and of the camera path, with an uncertainty in the certificate. The older T3 target (QR) is still recognised, but only up close.
3. One target per scan. The app tells fiducials A–D apart by their QR content; two targets in one scan would mix in the metrology. If you take an overview and details, each scan has its own (the same) target, placed anew at the captured site every time.
4. The target must be captured close up at least once. A QR code can only be read when it has enough pixels in the image. A 40 mm target is about 86 px at 1.2 m, which is only enough for a perfectly sharp frame — the slightest motion blur ruins it (verified by measurement). So during the lap slow down once and pass over the target from 0.5–0.8 m; it is then 130–200 px and decodes reliably even with slight movement. Watch the target counter on screen — while it shows zero, the target has not been captured and the scale will not be documented.
5. Placement: flat, in the plane of the captured surface, fully in view on at least 10 frames, not at the edge of the image. Body overview: 20–30 cm from the side of the trunk (waist level), on the surface, not on the body. Detail of a finding: 5–10 cm from the finding, in the plane of the skin (on the body or on the surface right beside it) so it is at the same distance from the camera as the finding. Object: beside it on the same base. The target must not be bent or wavy — glue it to a rigid board.
6. Light: even, diffuse (ceiling light, overcast sky). Turn the flash off. Direct sun interferes with LiDAR; hard shadows spoil the texture.
7. Surface: dry glossy and wet areas if possible; transparent and uniform surfaces are hard for photogrammetry — LiDAR holds the geometry there.
8. Space: walk freely around the subject. If you cannot (body against a wall), the hidden side will be missing in the model — note it in the remarks.
9. Diagrams: top view (two laps, target), side view (heights and tilt) and a detail of a finding are in the figures below.
10. Sound: the app signals progress by sound — theme Deep Space (default: a quiet deep tone and an airy whistle, unobtrusive) or Star Trek (bridge sensor warble); if it is not appropriate on site, choose No sound with the speaker icon on the capture screen. The pitch rises with coverage; a new sector and completion have their own signal.

Overview of a lying body — two laps around the body, one T3 target

one target per scan

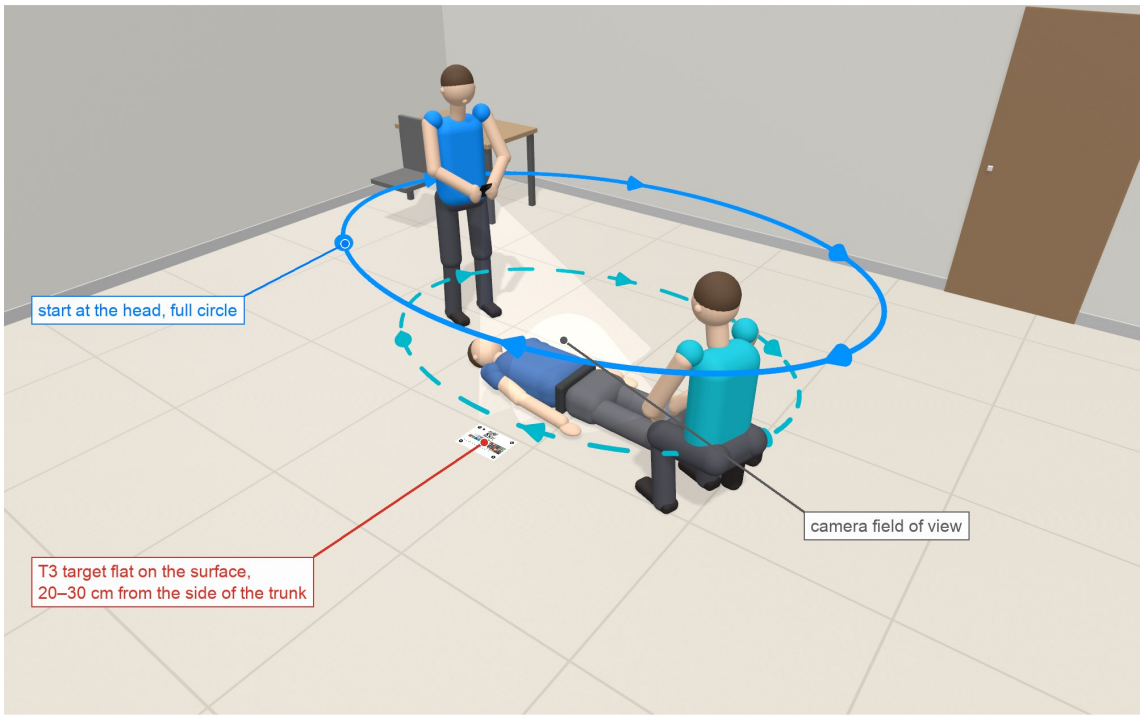

Lap 1: standing, phone at waist height, tilted 45° down, 0.7–0.9 m from the body, one step per second
Lap 2: crouching, phone 30–50 cm above the surface, almost level, 0.5–0.7 m from the body

Overview of a lying body from above: lap 1 from waist height tilted 45° (0.7–0.9 m), lap 2 crouching almost level (0.5–0.7 m); one T3 target flat on the surface 20–30 cm from the side of the trunk.

Side view — phone height and tilt

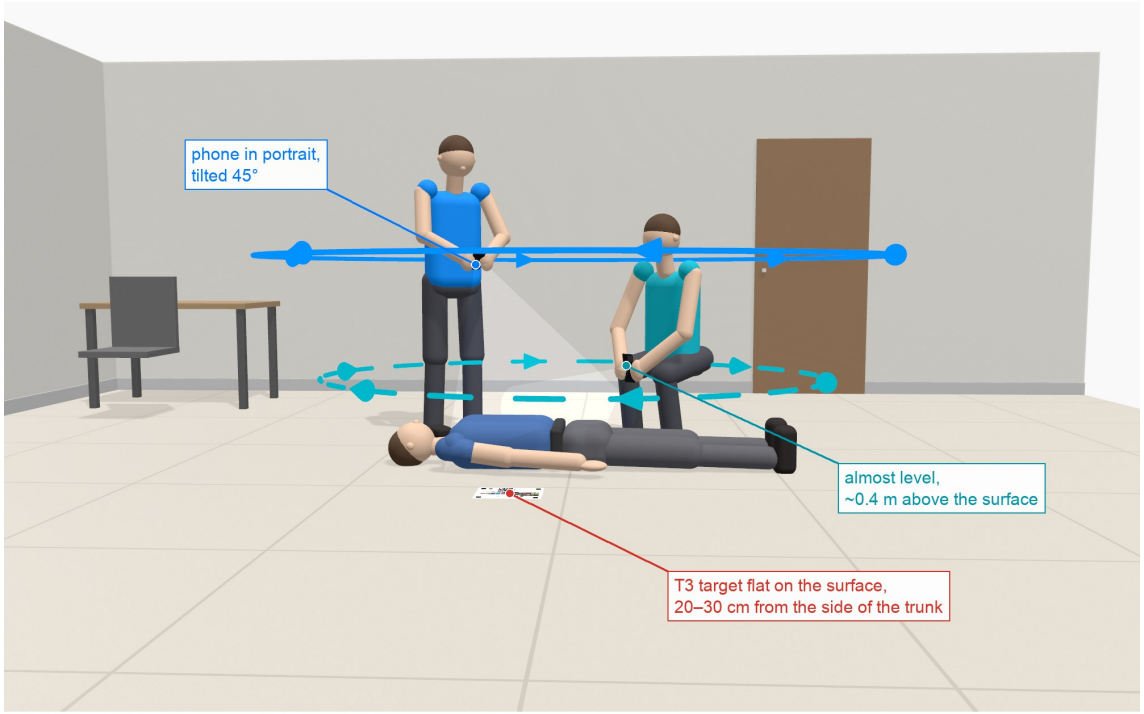

Lap 1: standing, phone at waist height, tilted 45° down, 0.7–0.9 m from the body, one step per second
Lap 2: crouching, phone 30–50 cm above the surface, almost level, 0.5–0.7 m from the body

Side view: heights and tilt of the phone for both laps; target in the plane of the surface.

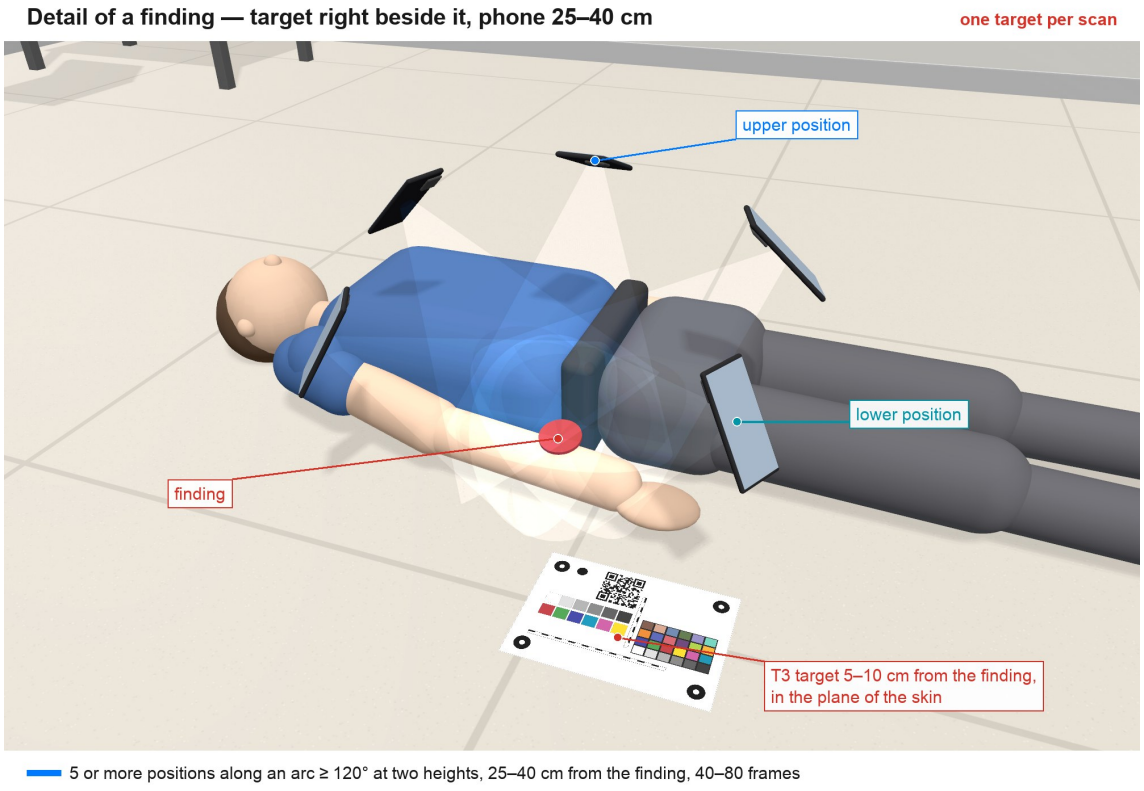

Detail of a finding: target 5–10 cm from the finding in the plane of the skin, phone 25–40 cm, an arc of at least 120° at two heights, 40–80 frames.

3. New scan — what the options mean

1. Tap +. Enter the case reference, examiner and place — they go into the report and the package name.
2. Profile: Deceased — examination (densest capture, 200 frames, accuracy) · Living person (quick capture, 120 frames, the person must not move) · Object / item (stone, tool, trace; circle at three heights).
3. Extent: Overview — whole subject, or Detail — region (wound, tattoo, hand). The triangle budget is the same, so Detail gives a much finer model. Whole body and injury detail are two separate scans.
4. Mode: Silent — 4K video (recommended: no shutter sound, depth exactly synchronous with the image) or Photo — full resolution (12 MP frames, slower pace).
5. Also save surroundings from LiDAR (on by default): the room mesh is collected during capture — position of the body in the room, distances from walls and objects.
6. Press Start capture.

3.1 Video or photo — when and why

Silent — 4K video takes frames straight from the ARKit stream (3840×2160, cropped to the LiDAR aspect ≈ 6 MP): depth comes from the same instant as the image, no shutter sound, a frame every ~6 cm of movement, low memory and thermal load. Photo — full resolution calls the system camera (12 MP, photogrammetry limit 4096 px): depth is paired with the nearest LiDAR measurement in time (within 0.2 s — with fast movement the frame stays without depth), the shutter is audible according to the ring switch, one frame takes ~0.6 s, the phone heats up sooner.

- Video: geometry and scale (exactly synchronous depth), smooth capture of 200 frames in 2–3 minutes, tolerant of breathing and micro-movement, silent. For overviews, scenes, living persons and everything where shape and dimensions matter.
- Photo: clearly better texture (twice the resolution), but slow, sensitive to movement (blur, depth from another pose) and thermally demanding at 200 frames. For static details where texture matters: wounds, tattoos, abrasions, small objects, blood stains.
- Phone orientation: frames are stored in the sensor's native orientation and gravity goes with every frame, so the UI orientation has no effect on the model. On iPhone the capture screen is locked to portrait (from 2.2.1) so it does not re-layout while you walk; hold the phone as is comfortable — portrait for a body (more overlap along the body), turned when sweeping room walls. Do not change the grip abruptly during one scan.

Rule: geometry and scale → video, texture → photo. Blurred frames ruin a measurement more than lower resolution does.

3.2 Recommended settings by situation

1. Deceased body — examination: profile Deceased, extent Overview, video, T3 target beside the trunk, surroundings from LiDAR on; three heights, 0.8–1.2 m, 150–200 frames; reconstruction extent Subject. Every finding you will measure: a separate Detail scan in Photo mode, 0.3–0.5 m, 40–80 frames, target right next to it — the body does not move, so photo gives the most here.
2. Body and surroundings (homicide, suicide, blood stains, objects, position in the room): the whole first, details after, move nothing. Scan A: Deceased, Overview, video, surroundings from LiDAR on — circle the body, then slowly sweep the camera over the floor, walls, furniture and objects within

2–3 m; reconstruction extent Subject with surroundings (body + 1.5 m), Whole scene only if you need distant parts of the room photogrammetrically too (the triangle budget is fixed, a whole-room mesh will be coarse). Measure distances of the body from walls, doors and objects on the Surroundings from LiDAR — metric from the sensor, independent of photogrammetry, with its own hash in the report. Scans B, C…: every blood stain, weapon, trace on the floor as profile Object, extent Detail, Photo, target beside it in the plane of the trace; blood stains are glossy — diffuse light, no flash; capture spatter on a wall obliquely from three directions, not head-on. Note in scan A which detail scans belong to it.

3 Living person: profile Living person, Overview, video (under a minute, breathing does not matter, silent); the person stands or lies still, arms slightly away from the body, gaze fixed. A specific injury: Detail, again video, 30–40 cm, 40–60 frames, target in the plane of the skin. Photo only if the person can hold still for 1–2 minutes (a resting limb, a supported back) — with movement photo is worse than video. Neck, face and hair are hard for photogrammetry: closer and slower.

4 Object (stone, knife, tool, shoe): profile Object, Detail, Photo, matte contrasting background, target beside it, three heights, 100–200 frames; the underside as a second scan after turning. Large object (car, furniture): Object, Overview, video.

3.3 Extent: Overview or Detail

iOS gives the mesh a single detail level, so the triangle budget is the same for a whole body and for a wound. Overview (radius ~1.2 m) spreads it over the whole body — position, body length, coarse measurements. Detail (radius ~0.35 m) gives the same density to a 30–50 cm region: a 5–10× finer surface, reliable millimetre measurements. Hence: one body = one Overview + as many Details as there are findings to measure. A Detail larger than ~60 cm adds nothing — use Overview with reconstruction extent Subject instead. Reconstruction extent (after capture, can be changed and recomputed): Subject = crop to the volume around the centre, the whole budget goes to the body; Subject with surroundings = +1.5 m around for context; Whole scene = everything, coarser; the whole scene remains as a second viewer even when cropped.

4. Capture step by step

1 Aim at the centre of the subject (for a body at the trunk) and tap Set scan centre. Coverage is computed around this point.

2 Walk slowly at about 0.8–1.2 m (object 0.3–0.5 m). Frames trigger themselves by translation and rotation of the camera — you press nothing.

3 Watch the radar top right: 36 sectors (10° each) × 3 height rings (outer = from above, middle = level, inner = from below). White wedge = where you stand; coloured = covered. Coverage 100 % = 30 of 36 directions covered at any height (Deceased; Living person 24, Object 33) — then the radar turns green. The rings need not be fully green, but for measurement two of them should be filled: middle and outer. The sound rises with coverage, a new sector adds a spark, completion plays a chord.

4 Cover two to three heights. With a lying body (most examinations): first lap from above — phone at waist to chest height, tilted down about 45° (outer ring); second lap low — crouch, phone 30–50 cm above the surface, almost level (middle ring). A third lap from above with an outstretched arm adds the crown, face and the upper side of the hands. Every point of the surface should appear on at least three frames from different angles.

5 Guidance (from 2.2.1): the bar above the Finish button says what to do now — Slower (fast movement), Closer / Further away (distance from the surface outside the recommendation), Aim at the centre (camera off the subject), Step back a little (an uncovered sector behind you), Second lap lower / Next lap higher (one ring done, the other empty), Low light, Coverage is sufficient. An orange hint comes with a light buzz and a tone, green means done. The hint changes at most 4× per second, so you have time to react.

6 If “Moving too fast” or “Few details in view” appears, stop, wait for recovery and continue more slowly.

7 Tap Finish capture (at least 20 frames, ideally once the radar is green). Processing starts by itself.

4.1 Deceased body — examination

- Target beside the trunk, visible from all sides. The body lies: start at the head, circle the whole body from above (phone at waist height, tilted 45° down, 0.7–0.9 m from the body, one step per second), then a second lap low (crouch, phone almost level, 0.5–0.7 m) — this gives the sides of the trunk, the sides of the limbs and the profile of the face; without it the model looks “smeared” from the side. A third lap from above with an arm stretched over the body adds the crown and the upper side of the hands. At a wall circle what you can — 30 directions are enough.
- Slow down at the head, hands and feet and add close frames (30–40 cm) — the most frequent sites of findings. The app guides you with hints — follow them; when Coverage is sufficient shows and both radar rings are filled, finish.
- Injuries you know you will measure: capture additionally as a separate scan with extent Detail and the target right next to them.

4.2 A living person with the T3 target — the full procedure

In one respect a living person is an easier subject than a body: you can tell them how to stand. In another they are harder — they move. The whole capture is therefore a race against time: prepare everything first, then walk three laps without stopping. The Living person profile is built for this — it needs fewer frames and tolerates slight movement, so the capture can be done in under a minute.

1. Explain the procedure. Say it will take about a minute, that you will walk around them and that they must not move. Handle consent and instruction according to your own rules — the app sends nothing anywhere, the data stay on the phone.
2. Place the person in an open space. At least 1.5 m of free space all around so you can walk a full circle. Even light, never against a window — backlight turns the figure into a silhouette without detail.
3. Put the T3 target on the floor beside the feet, about 20–40 cm to the side, completely flat and not covered. One target per scan. It need not be in every frame, part of the lap is enough.
4. New scan: profile Living person, extent Overview of the whole body, target on. Aim at the chest and tap Set scan centre — coverage and the crop are both computed around this point.
5. Walk three laps at a distance of 1.0–1.2 m (see the figures below): first at head height tilted down, then at chest height level, finally crouching with the phone tilted slightly up at the legs. One lap takes 15–20 seconds, at a steady pace, without stopping.
6. Finish when the radar turns green and at least two rings are filled. Processing starts by itself.

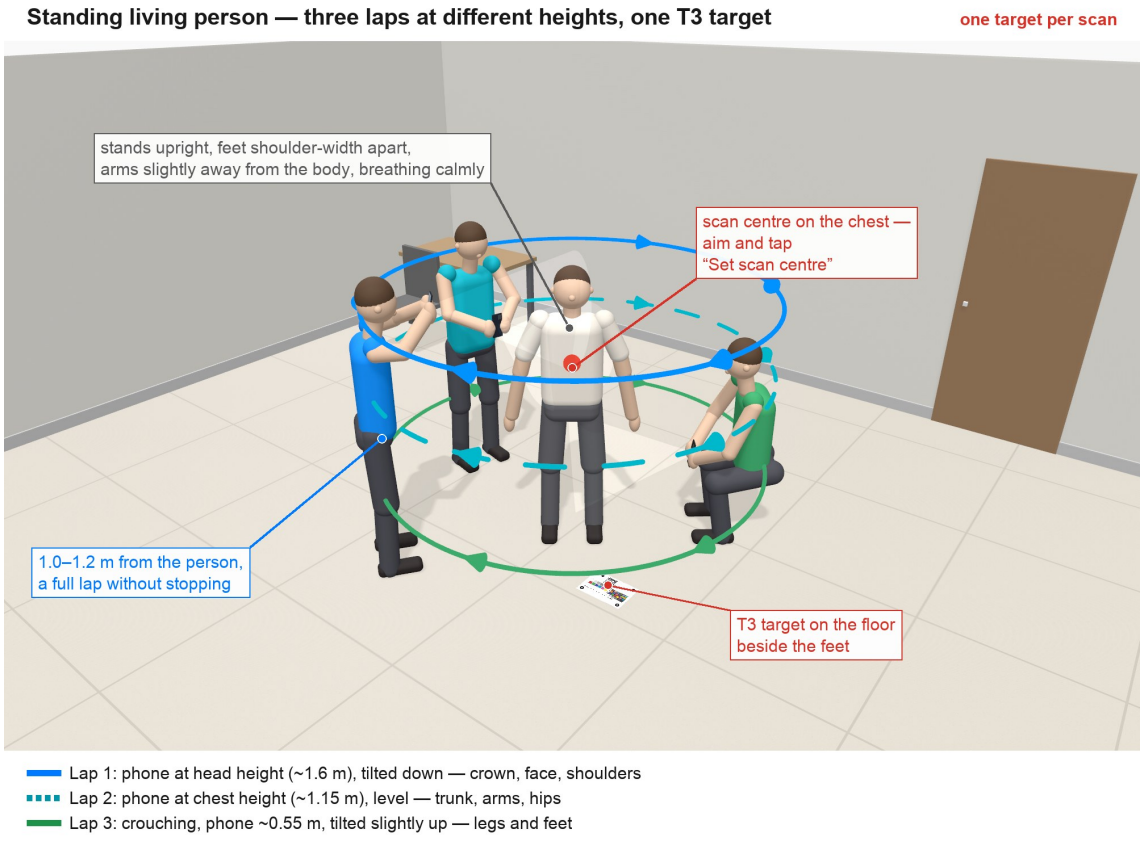

Standing person from above: three laps at different heights on a circle of 1.0–1.2 m, the T3 target flat on the floor beside the feet, the scan centre on the chest.

Living person — three phone heights above each other

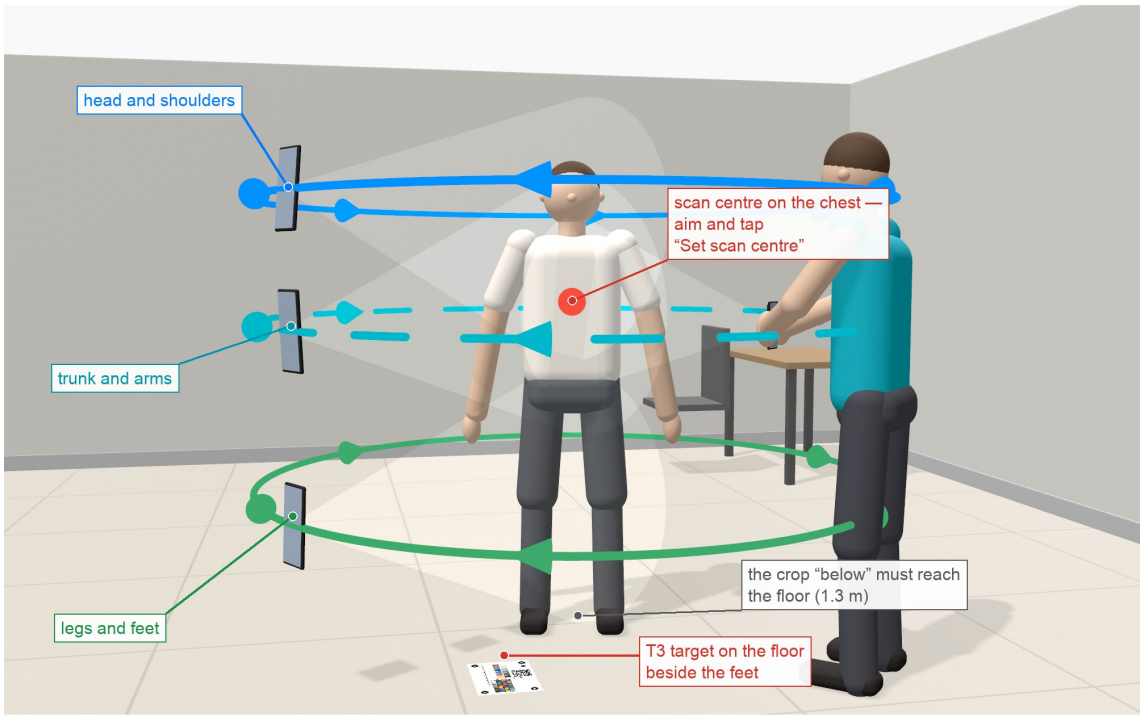

The same three laps from the side: phone at head height (tilted down), at chest height (level) and crouching (tilted slightly up). Together they cover the crown, the trunk and the feet.

4.2.1 Posture — what matters most

- Arms slightly away from the body, palms to the side or forward. Arms pressed against the sides are the most common mistake: the inner surfaces of the arms and the sides of the trunk will then be missing from the model, because the camera never saw them.
- Feet shoulder-width apart, weight on both. The inner surfaces of the thighs then get both light and coverage.
- Head straight, gaze on a fixed point ahead. Turning the head to follow the operator is the second most common mistake — the head then comes out blurred or doubled in the model.
- Breathe calmly. Chest breathing does not bother photogrammetry; shifting the feet, swaying and turning do.
- If the person cannot stand: seat them on a chair without a back in the middle of the free space, put the scan centre on the chest and do the third lap (legs) crouching from the front and the side. A backrest and armrests hide the back — document the back with a separate scan in that case.

4.2.2 Clothing and skin

- Whatever is to be documented must be uncovered. Nothing can be measured through clothing.
- Smooth single-colour clothing is the worst case for photogrammetry (a white T-shirt, black trousers) — there is nothing for the points to attach to and the surface comes out as a smeared bubble. Skin has pores, hairs and pigment, so it works far better. If a plain surface cannot be avoided, use sensitivity High when processing.

- Glossy surfaces (oily skin, latex, a shiny blouse) reflect in a different place in every frame and the algorithm cannot match them. Oily skin can be gently dried.
- Hair is never reconstructed faithfully — for photogrammetry it is transparent. If a finding is in the hairy part of the head, make a separate Detail scan of it with a target.

4.2.3 After the capture

- Crop: a standing person has the scan centre on the chest, about 1.3 m above the floor — the below value must therefore be 1.3 m, otherwise the legs are cut off. From version 2.2.1 the app sets this by itself for the Living person profile, including older scans; you only need to change it for a seated person (0.8 m is enough) — see chapter 5.3.
- Findings to be measured (a wound, a bruise, a scar) should also be captured as a separate scan with extent Detail and a target 5–10 cm from the finding in the plane of the skin — a whole-body overview does not have enough triangles for millimetre measurement.
- The face is personal data. If you are documenting only, say, a wound on the forearm, make a Detail scan of the forearm alone — the model then contains no face and there is nothing to protect.

If the person moved during a lap — shifted their feet, turned their head, lowered their arms — repeat the scan. Photogrammetry cannot process movement: the model will have doubled outlines or holes and a measurement taken from it will not be valid. Capturing a second time is better than measuring on a wrong model.

4.3 Object

- Matte, contrasting background (not a glossy table), target next to the object. Circle at three heights: from below, level, from above. Capture small objects close up with extent Detail.
- Document the underside with a second scan after turning the object over.

4.4 Scene / room

- After circling the body, sweep the camera slowly over the walls, floor and nearby objects — the “surroundings …” counter (number of triangles) on screen grows. This fills the room mesh (chapter 9).
- For processing choose extent Subject with surroundings or Whole scene if the photogrammetric model should include the surroundings.

Most frequent capture mistakes: fast movement (blurred frames), skipped top height (crown, back of the hand), target out of frame or tilted, flash, capturing against a window. Each of them degrades measurement more than any processing setting.

5. Processing

Reconstruction runs entirely on the device, minutes (200 frames 10–20 min). The screen stays on and dimmed; leave the phone on the charger and undisturbed. A notification arrives on completion and the examination opens.

- Reconstruction extent: Subject (crop to the volume around the centre — the whole triangle budget goes into the body) · Subject with surroundings · Whole scene. The whole scene remains available as a second view.
- Sensitivity: High helps with pale skin and low-texture surfaces (takes longer).
- Recompute at any time without new capture — e.g. another extent or sensitivity. The app asks for confirmation: every computation gives the model its own coordinate frame, but findings and measurements are transferred to the new model by the camera positions (from 2.2.2; the original file is kept aside and the report states the transfer). The previous model is kept aside during the computation — on cancel or failure it is restored, so you are never left without a model.
- Processing runs on iPhone and iPad with A17 Pro / M1 or newer. The aPrehliadky app on a Mac cannot compute (chapter 5.2); a package with a finished model opens there for examination, measuring and the report.

5.1 The app quit during processing (E36) — what happened and what to do

Photogrammetry of 200 4K frames is the most demanding task an iPhone can run: halfway through it holds gigabytes in memory. iOS terminates the app without warning when (1) it goes to the background — locking the phone, switching to another app, a call; a suspended app with a large memory footprint is first in line, (2) memory runs out (little free RAM, many open apps), (3) the device overheats. Processing then ends silently, typically around 40–60 %. This is not a capture error — the frames are fine and stay stored.

1. When you reopen the scan the app itself shows E36 — processing was interrupted (with the attempt number, phase, percentage and whether it went to the background) and a button Start economy processing: frames are downsized to 2560 px (about 2.3× less memory, depth stays), extent Whole scene. Measurement accuracy is practically unchanged — scale comes from LiDAR, not from pixel count.
2. During processing keep the app in the foreground on the charger: the screen dims by itself (it does not turn off), do not lock the phone, do not switch to another app, close other apps (Safari, Maps, games). If you need the phone for something else, postpone the processing.
3. If it recurs even in economy mode: sensitivity Normal, extent Whole scene, let the device cool down. The processing log records free memory and temperature every 10 % — attach it when reporting.
4. Alternative: recompute the package with frames (Export → Model, findings and frames) on another device or an iPad with more memory; the model can be sent back as a package.

Lesson from a scan that ended at 50 %: 200 frames in 45 s (4.5 frames per second) is too fast — 54 of 108 sectors were covered and the camera was on average 0.7 m from the centre. Slower circling at three heights gives a better model and a lighter processing load; after finishing the capture keep the app in the foreground until the “Model is ready” notification arrives.

5.2 Computing on a Mac — older iPhone, higher detail

Any iPhone with iOS 18 can capture (even without LiDAR: scale then comes from the camera path and the T3 target), but only an iPhone/iPad with A17 Pro, M1 or newer can compute the 3D model. The solution is the native Fotograme3ka 3D app for Apple Silicon Macs (App Store, free): it recomputes the package with frames at medium, full or custom detail (polygon cap, texture up to 16K) — finer than the iPhone — and returns a package with the model. From 1.2.0 it is also a full examination tool: findings on the skin, measuring along the surface, lasso with vertex editing, report and certificate PDF, scale with uncertainty. Several packages dropped at once are processed as a batch (results saved next to the inputs as -mac.f3d).

1. On the iPhone: scan → Export → Model, findings and frames → AirDrop to the Mac (the .f3d package has hundreds of MB; frames are required).
2. On the Mac: open Fotograme3ka 3D Computation, drag the package into the window, choose detail (Medium recommended) and extent (Subject), press Compute model. Medium detail with 200 frames takes 7–15 minutes, full detail several times longer.
3. When finished, Save package… and AirDrop it back to the iPhone; opening the package imports it as a new scan with the model, the whole scene, LiDAR surroundings and the processing log. Issue the certificate anew on the iPhone (the old one belonged to a different model).
4. If findings were already entered, the Mac transfers them to the new model just like the iPhone does (chapter 5.3); the imported scan carries a transfer record and the estimated position error.

The aPrehliadky app for Mac (“Designed for iPad” from the App Store) is for viewing, measuring and reports; do not compute the model there — the iOS photogrammetry layer crashes on a Mac. Use Fotograme3ka 3D Computation for computing.

5.3 Practical examples and explanations

The crop is not the phone distance

The three numbers under Crop on the processing screen (radius, above, below) say nothing about how to hold the phone. They are the dimensions of an invisible box the app places around the scan centre — the point you aimed at during capture and tapped “Set scan centre”. Processing keeps only what is inside the box and discards everything outside. This is done so the whole triangle budget goes into the body and not into the wall behind it.

- radius — how many metres the box reaches sideways, forward and backward
- above — how many metres it reaches up from the scan centre
- below — how many metres it reaches down from the scan centre

What matters is where the scan centre sits. With a lying body it is on the trunk, about 20 cm above the floor, and 0.5 m down is plenty. With a standing person it is on the chest, about 1.3 m above the floor — a box reaching only 0.5 m down then ends above the knees and the legs never make it into the model.

The crop is not the phone distance. It is an invisible box around the scan centre.

The scan centre is the point you aimed at and tapped “Set scan centre”. Processing keeps only what is inside the box.

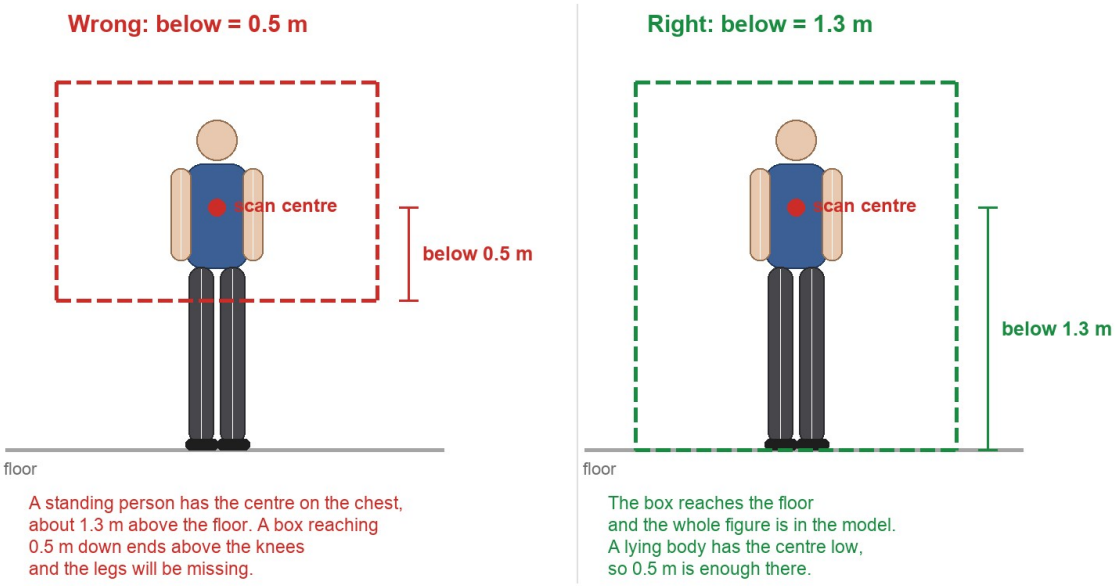

The same “below” value, two completely different results. What decides is the height of the scan centre above the floor.

From version 2.2.1 (build 32) the app corrects the old 0.5 m value to 1.3 m by itself for the Living person profile, including packages captured with an older version. You need to do nothing. If you do need to change the numbers — a seated person, a body on a high table — the sliders are on the processing screen.

The phone distance is something else

- Body overview: 0.7–0.9 m from the body, one step per second.
- Detail of a finding: 25–40 cm from the finding.
- Object: 0.3–0.5 m, a small object even closer.

This is not set anywhere in the app. You set the distance by where you stand, and the guidance bar tells you with Closer or Further away.

Why the scale row is orange when the model is fine

The row ARKit trajectory vs. model compares the camera path recorded by ARKit (LiDAR + IMU, metric) with the camera positions photogrammetry found. From version 3.0 this path is the source of the measurement scale: a photogrammetry model is not in metres by itself, even with LiDAR depth (the same scan came out 6 % smaller once and 14 % smaller another time), so the app converts every measurement with a coefficient from this path and the viewer shows the correction and its uncertainty. The second, independent source is the T3 target triangulated in the model; their agreement is the independent check for class A. An orange row means the fit is not precise enough for a correction (residual over 25 mm or fewer than 50 cameras) — the scale is then only indicative.

What happens to findings when you recompute

Every computation produces a model in its own coordinate frame, so old pins and measurements would sit in the wrong places on the new model. From version 2.2.2 the app transfers them: the same frames have their cameras at the same places in the old and the new model, and from these pairs the app computes how the new frame is rotated, shifted and scaled relative to the old one. The same transformation moves the pins and the measurement points; the error is usually a few millimetres and the app shows it (“estimated position error”). For older scans without stored camera poses the transfer goes through the ARKit camera path (coarser, about a centimetre), and if even that is not reliable the findings are only set aside and you enter them again.

- Measurement values do not change — they were measured on the original model and are physical (mm). The transfer moves only the points. If you want the value from the new model, measure again.
- The original findings file stays in the scan folder as annotations_<date>.json (it goes into the archive package too) — the chain of evidence is preserved.
- After a recompute the viewer shows a bar “Findings transferred…” and the findings list carries the transfer record; the report states the method, the number of cameras and the estimated error. Check the pins — after a large change of extent (Subject → Whole scene) the surface in the new model may lie slightly elsewhere.
- The previous model is kept aside during the computation — if you cancel it or it fails, it is restored. Issue the certificate anew, the old one belonged to a different model.

6. Examination — controls

- One finger rotates, two fingers pan, pinch zooms. View menu: front, back, left, right, top, bottom.
- The bottom panel has two modes: Measure and Mark finding. The switch says what a tap on the surface does. Nothing is hidden in gestures — a long press is only a shortcut to the full finding form.
- The list icon top right opens all findings and measurements; the palette icon turns the texture off (bare geometry shows holes and artefacts); the share icon opens export. The … menu has a toggle Flip model upside down. Apple photogrammetry takes the gravity direction solely from HEIC metadata (Apple MakerNote); frames from versions before 2.2.1 (build 26) lacked it and the model came out inverted. From build 26 gravity is written into the frames themselves, a recomputation of an older session adds it automatically, and the trajectory fit recognises an inverted model so the viewer flips it by itself; the toggle changes only the display, not the data.
- AR 1 : 1 (ARKit icon, iPhone/iPad only): place the model on a table at real size and compare with a real ruler.

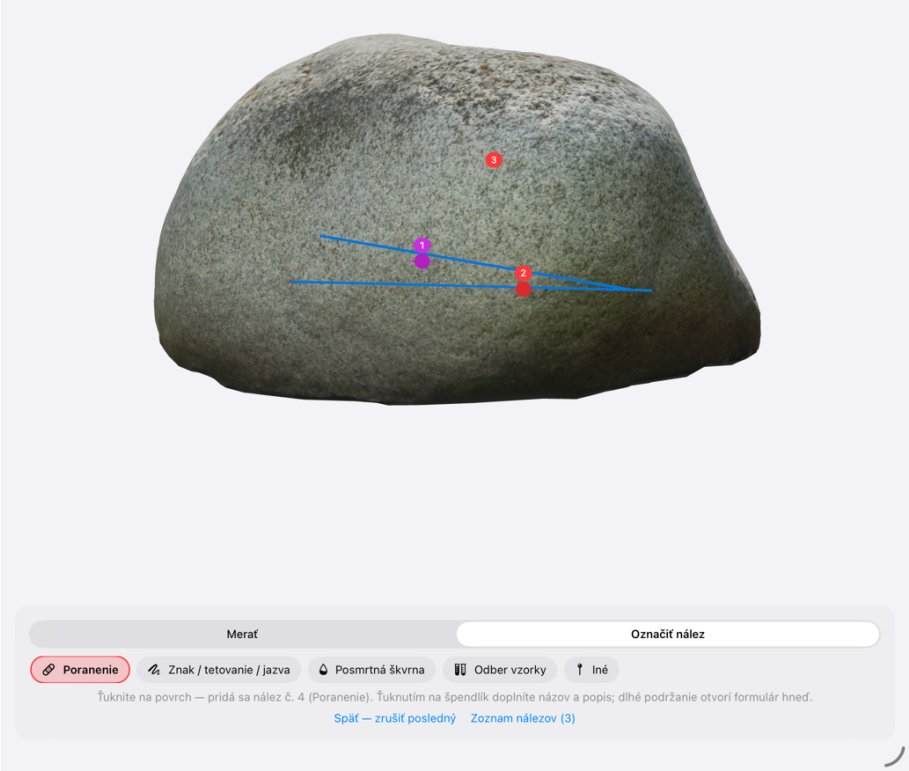

Mark finding mode: choose the type once, every tap on the surface adds a numbered pin (here no. 3).

7. Findings — how to mark precisely and describe

1 Switch the panel to Mark finding and choose the type: Injury · Mark / tattoo / scar · Livor mortis · Sample · Other.

2 Zoom in so the finding fills half the screen and tap exactly its centre (for a wound the centre of the wound, for a bruise the centre of the area). The pin gets a sequence number and the type.

3 Tap the pin → Title (short, as in the autopsy report: “laceration right forehead”) and Note (location relative to landmarks, dimensions from measurement, colour, margins, depth, direction).

4 Mistake: Undo — remove the last one deletes the last pin. A misplaced older finding: tap it → Move on the model → tap the correct place.

5 Numbering is preserved; in the report findings are listed in order with their position in the model (x, y, z in metres), so they can be found again.

Mark while zoomed in, never from afar. One finding = one pin at its centre; dimensions belong to measurement (chapter 8), not to an estimate. Title = what and where, note = how it looks. Use the same words as in the autopsy report — a colleague must pair them without explanation.

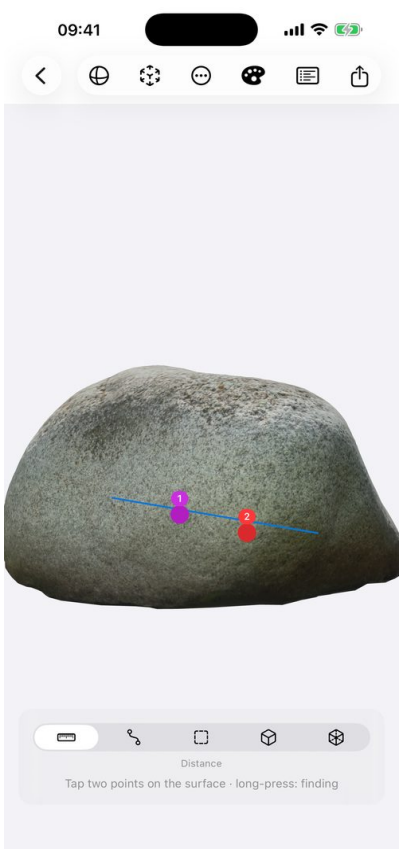

Stone — examination with two findings (labels 1 and 2)

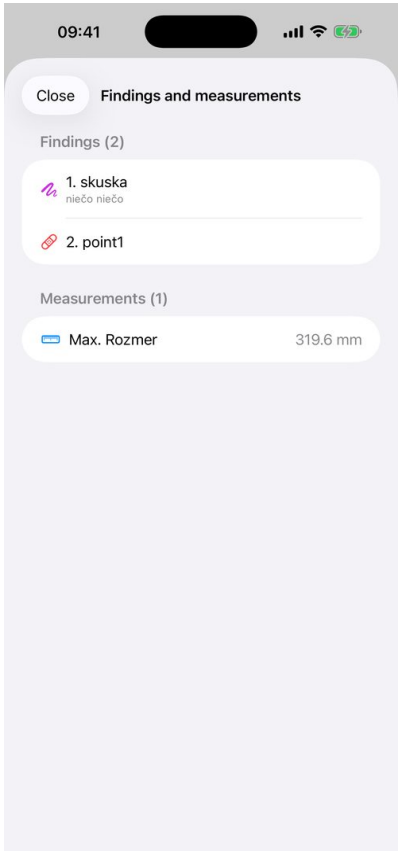

List of findings and saved measurements

Canister — Object / item profile, Detail extent

8. Measuring with the smallest error

Switch the panel to Measure and choose the kind: Distance (2 points, straight line) · Path (along the surface — wound length on a curved body, circumference) · Area (polygon ≥ 3 points — bruise, abrasion) · Volume — box (bounding box from points) · Volume — model (whole mesh, indicative).

1. Reference first, everything else after. Measure the 100 mm scale of the target with Distance (line to line, zoomed in), tap Reference, enter 100. The deviation is stored in the report; it should be within ±1 %. If larger, repeat the scan with the target closer to the subject.
2. Zoom in so both ends of the measurement span the screen. Tap edges and clearly defined points (end of the wound, margin of the bruise), not “somewhere near them”.
3. Use Distance for dimensions across a flat or slightly convex surface (body length: crown–heel). If “along the surface” also appears under the value, the surface bulges between the points — record both values and say which is which.
4. Use Path for wound length on a curved surface and for circumferences: tap points along the line every 2–5 cm; between taps the length adds up along the real mesh surface. A path is never shorter than the straight line.
5. Tap Area around the outline of the finding with 6–12 points in order (not criss-cross). Area is a planar projection — on a strongly curved surface (shoulder, calf) split it into two smaller ones and add them.
6. Volume is indicative: the box from points gives an upper bound, the model volume depends on mesh closure and surroundings. Report it as indicative.
7. Save right after every measurement; the label may stay empty (“Distance 3” is filled in) or match the finding (“wound length no. 1”).
8. Repeat key measurements twice from another view; if they differ by more than 2 mm on a detail or 1 cm on a whole body, take the decisive measurement on a scan with extent Detail.
9. Units ((i) → Settings): mm, cm, dm, m, inches, feet; the conversion is exact, the package always stores mm.

A measurement is defensible when: (1) the scale comes from the camera path (green row) and is substantiated by the T3 target in the model or a reference within 1 %, (2) points were placed zoomed in on clear edges — with the loupe (press and drag), an Apple Pencil, or a mouse on the Mac, (3) on a curved surface you used Path, not Distance, (4) the record states the value with its uncertainty: “aPrehliadky 3.0 3D model, 123.4 ± 1.8 mm (k = 2), class A”. A 0.1 mm display resolution is not accuracy — the uncertainty U determines accuracy.

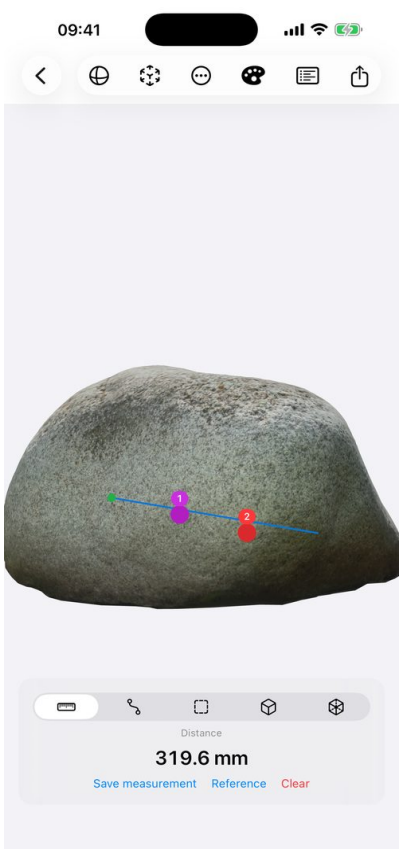

Distance: two points on the surface

Area: four corners of the label (example)

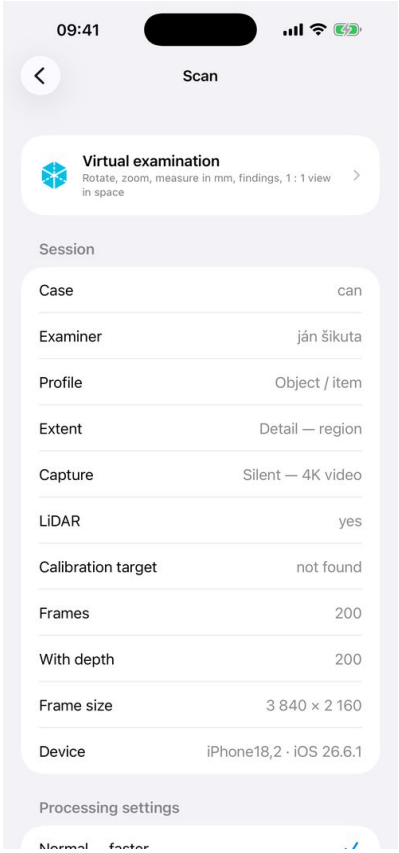

Scan screen: parameters, processing, diagnostics

9. Surroundings from LiDAR — room mesh

With LiDAR the app also collects the ARKit room mesh during capture. After capture a second view appears in the scan, Surroundings from LiDAR: a metric, untextured model coloured by classification (floor, walls, ceiling, table, seat, window, door). It is a direct LiDAR measurement, not photogrammetry and nothing is filled in.

- Suited to the position of the body in the room and distances from walls, doors and furniture; it has its own findings and measurements (different coordinates from the body model) and its own report section with SHA-256.
- Mesh resolution is in the order of centimetres — measure small wounds on the photogrammetric model, the distance of the body from a wall on the mesh.

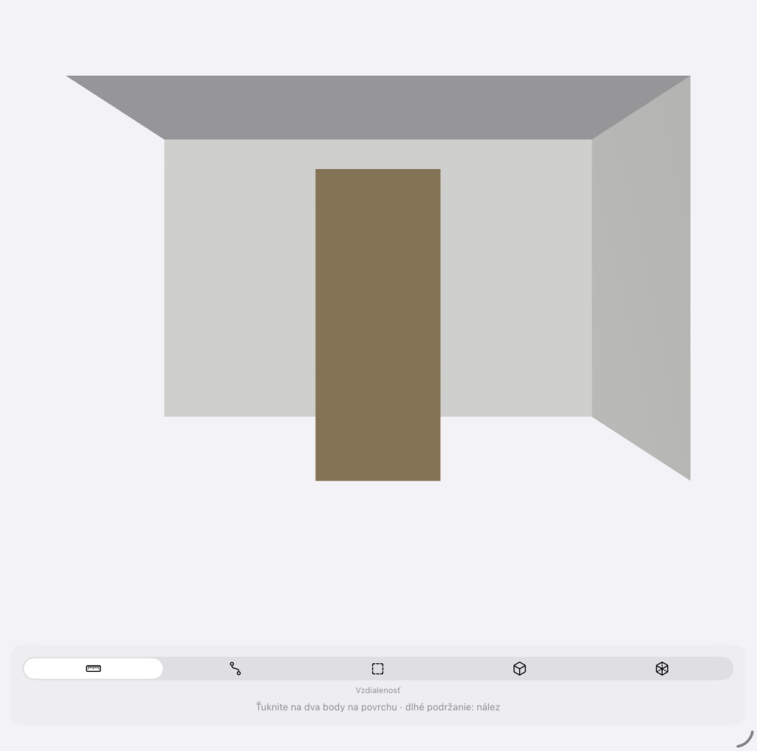

Surroundings from LiDAR: colours by classification (walls, ceiling, door), measurement in metres.

10. Scale — when a measurement is binding

Photogrammetry alone does not know scale and — as the 3.0 audit showed — does not keep it even with LiDAR depth: the same scan came out 6 % smaller once and 14 % smaller another time. aPrehliadky 3.0 therefore always derives the measurement scale from independent sources, fuses them and states in the report which applies and with what uncertainty:

1. LiDAR depth: every frame carries a depth map embedded in the HEIC — photogrammetry uses it as a prior (shape, rough scale), but the app does not take the measurement scale from it. Report line: “HEIC carries embedded LiDAR depth”.
2. Camera trajectory (ARKit): camera positions are metric (LiDAR + IMU). After processing they are matched with the photogrammetry poses and yield the measurement scale (from 3.0 always, even with LiDAR — a photogrammetry model is not in metres by itself) with an uncertainty (typically ~0.5 %). The viewer shows the correction (“Scale corrected by x %”) and every measurement carries an expanded uncertainty ± U (k = 2). The second independent source is the T3 target triangulated in the model (uncertainty ~0.1–0.3 %); both are fused and their difference goes into the certificate.
3. T3 target in the model: the app finds the circular fiducials in every frame and triangulates them directly in the model after processing (10–15 known distances 33–294 mm) — a scale independent of LiDAR and of the camera path, uncertainty ~0.1–0.3 %. With two to four targets around the body enter the Target control network on the scan screen: distances between target centres measured with a steel tape — a 1–2 m baseline gives a global scale under 0.1 %.
4. Reference: the only check a court or a colleague can verify by eye on the model (the target's mm scale legible in the texture, a ruler). The deviation in % goes into the report.

Class A: scale from the camera path (green row) and an independent check by the target in the model or a reference within 1 % — measurements binding within the uncertainty U. Class B: metric scale without a second proof — state “no reference” in the report. Class C: scale only approximate (no fit, or an older scan without correction) — indicative; add a reference afterwards by measuring an object of known length in the model.

11. Report, export, archive, Mac

- PDF report: case, examiner, capture parameters, method statement (no NeRF/3DGS/AI), scale verification, findings with positions, measurements in the chosen units, surroundings from LiDAR, integrity, SHA-256 of model and mesh, limitations of the method.
- .f3d package (Fotograme3ka 3D; older .aprehliadka files import the same way): manifest, model, findings, certificate, room mesh, processing log, SHA-256 of every file; optionally the frames (allow recomputation elsewhere). The recipient sees on import whether the package was modified. With the Pro layer the package is additionally signed with the device key (Secure Enclave): the report states “Package signature: VALID, key …, signed by …” and the institution name in the header ((i) → Pro).
- 3D documentation certificate: Export → Issue certificate (for a finished model). It states the documentation class A–D (A = scale substantiated by two independent routes within 1 %, B = metric without a second proof, C = indicative, D = unverified), the capture trace (camera path, coverage, time), scale verification (LiDAR, trajectory fit, reference, T3 target), the SHA-256 of the model and findings, a signature (Pro) and a QR code. It is also stored in the package (certificate.json) and import verifies it. Methodology: web → Certification methodology (PDF).
- Archive: full package into Files → aPrehliadky → Exporty → Archív; the package SHA-256 is written into the manifest and the report. Frames can then be deleted from the phone.
- Save to… on export opens the system folder picker (iCloud Drive, external drive); on the iPhone exports are also in Files → On My iPhone → aPrehliadky → Exporty.
- Mac (Apple Silicon): double-click a package from the iPhone — examination, measuring, findings and the report work; capture and processing stay on iPhone/iPad.

11.1 What to archive for court and the investigator

Keep the archive package with frames in the case file — in the app Export → Archive. It is the complete, recomputable original: frames (raw data), model, whole scene, LiDAR surroundings, findings, measurements, certificate, processing log and a SHA-256 of every file. In this chain the frames are what negatives are in photography: without them the model cannot be verified independently, with them anyone can recompute and compare it. The SHA-256 of the whole package is written into the manifest and the report, so it can be proven exactly what went into the file.

1. Archive package …-archiv.f3d (hundreds of MB) — the main backup. Store it on two independent media; record its name, size and SHA-256 from the report.
2. Report PDF — what goes into the paper file and what the court reads.
3. 3D documentation certificate (PDF) — class A–D and the substantiation of scale. Attach it to the report.
4. Package without frames (tens of MB) — a working copy for the investigator or an expert to open and measure on an iPhone, iPad or Mac. It is not a substitute for the archive.

After archiving you can delete the frames from the phone (an option during archiving) — the model, findings and report remain and the phone does not hold gigabytes. Recomputation is then no longer possible from the phone, but it is possible from the archive package.

Do not send the model.usdz file alone as the only evidence. It is only the result, without the manifest, hashes, certificate and raw data — it cannot show how it was produced or that it was not modified.

12. ePrehliadky (ÚDZS) and legal notice

aPrehliadky is designed as a possible add-on to ePrehliadky — the software of the Slovak Health Care Surveillance Authority for electronic post-mortem examinations. The .f3d package, the PDF report and the certificate are attached to the record; the link aprehliadky://scan/<id> leads back to the scan. The package format is open (docs/FORMAT.md) and usable in any other system.

aPrehliadky is an independent app by the author (© 2026 Ján Šikuta, MD, PhD). It is not a product of ÚDZS, and ÚDZS neither publishes nor endorses it.

13. Troubleshooting

- The device cannot compute the model: capture, Export → Model, findings and frames, send the package by AirDrop to a supported device, import and start reconstruction.
- I cannot see the model: Scan → Diagnostics → Open in the system viewer. If the model exists it always shows. If not, send the processing log (it is in the package) to forensika.eu@icloud.com.
- Target not found: print at 100 %, target flat, in as many frames as possible, no glare.
- Measurement disagrees with a ruler: verify the reference; if the deviation exceeds 1 %, repeat the scan with the target closer to the subject and a slower walk.
- The phone heats up: normal during processing; leave it on the charger and undisturbed.
- Language: (i) → Settings → App language (takes effect after restart).

14. Error codes — meaning and remedy

Every error in the app is shown with a code E-xx, a short meaning and a remedy; the original system message (domain and number) is in square brackets and should be quoted when reporting. The most common one is E30 (“system code 1” from photogrammetry).

- E01 Device does not support capture (ARKit). — Capture on an iPhone/iPad with iOS 18; on a Mac and in the Simulator only import works.
- E02 Frames or depth could not be written. — Free up storage, repeat the capture; if it recurs, restart the device.
- E03 Session without frames. — Capture at least 20 frames before finishing.
- E04 Device cannot compute the model. — Recompute the package with frames on an iPhone/iPad with A17 Pro / M1 or newer.
- E05 The .f3d package is damaged. — Request the package again (AirDrop, not a Mail preview); do not unzip and re-zip.
- E10 Processing cancelled. — Tap Recompute.
- E20 Storage full · E21 File not found (frames deleted after archiving) · E22 No access to the file (open via Files → Share → aPrehliadky).
- E30 Photogrammetry could not process the frames (system code 1 — frames could not be matched/aligned). — 1) Recompute with sensitivity High and extent Whole scene; 2) input Folder without depth (diagnostics); 3) repeat the scan more slowly, with more overlap (three heights, 0.8–1.2 m), no subject movement, no glossy or plain surfaces. Causes: blurred frames, little texture, subject movement, changing light, frames against a window.
- E31 Photogrammetry rejected the input frames. — Try Folder without depth; if it passes, the depth data is damaged (capture again in silent mode).
- E32 Output cannot be written · E33 Not enough space for processing (states how much is needed). — Free up storage: archive old scans and delete their frames.
- E34 Photogrammetry input is not valid (code 2). — Frame diagnostics, recompute with Sample sequence.
- E35 Processing interrupted by the system (code 3/4 — memory/temperature). — Cool down on the charger, close other apps, extent Subject.
- E39 Other photogrammetry failure (system code given). — Sensitivity High; if it recurs, send the processing log.
- E40 Unsupported ARKit configuration (Silent mode, disable surroundings) · E41 Camera/LiDAR unavailable (close other camera apps) · E42 Camera not allowed (Settings → aPrehliadky → Camera) · E43 Tracking failed (stop, textured surface, more light) · E44 Full-resolution frame failed (Silent mode) · E45 Too little detail in view · E46 ARKit write failed (storage) · E49 Other ARKit error (restart the app).
- E60 App Store (donor/Pro) — connection and sign-in; the fee is not required for any feature.
- E99 Unknown error — the domain and number are in brackets; send the processing log (Scan → Processing log → Copy) to forensika.eu@icloud.com.

How to report an error: the code (E30), the text in brackets (e.g. “com.apple.CorePhotogrammetry 1”), iPhone model and iOS version, app version ((i)), and the processing log. Do not send frames anywhere —

the log is enough.

15. Privacy and donor

The app has no network layer. Frames, models, findings and reports stay on the device until you export them yourself. No analytics, no advertising.

aPrehliadky is free and fully functional with no limits — capture, measuring, findings and the report are and will remain free. Anyone who wishes can become a donor ((i) → Donor). The Pro layer (package signing, institution on the report) is optional and is recognised automatically for donors.