Colorime3ka

Colour measurement of skin findings for forensic and dermatological documentation

Professional manual

App version 3.0

Ján Šikuta, forensic pathologist

a forensika.eu project

September 2026

Contents

1. What the app is — and what it is not

> Colorime3ka measures the colour of a skin finding from a photograph or the live camera, converts it to CIELAB and names it by the nearest reference colour. Beyond colour it recognises the morphological pattern by chromophore — that is, by which substance causes the colour.

It is an **aid to description and documentation**. It replaces the sentence "a blue-purple macule" with an objective number that is repeatable and can be entered into a report. Nothing more — and that "nothing more" is deliberate.

> **What the app is not.** It is not a medical device. It does not diagnose, does not screen, and never states that a lesion is or is not malignant. It does not determine the age of a bruise. It does not replace an examination or an autopsy. Percentages express the agreement of measured features with reference criteria — not a probability of disease and not diagnostic certainty.

The app uses no neural network and no trained model, and no image is ever analysed for disease. All outputs come from deterministic arithmetic following published formulae (sRGB → CIELAB, ΔE CIEDE2000) combined with fixed, human-authored threshold tables. It makes no network request of any kind.

2. Quick start

1. **A photograph** — pick one in the *Photo* tab, or measure from the live view in the *Camera* tab.
2. **Tap the finding.** The measuring area appears: the solid circle is the measured region, the dotted ring is the surrounding skin against which everything is computed.
3. **Set the size** with the slider so that the ring lies on healthy skin. The app advises you: "*the lesion is correctly framed*" is the target state.
4. **Choose the mode** with the shield icon, top right: colour only · forensic assessment · dermatological assessment.
5. **Open the assessment** by tapping the bottom panel, add the clinical context and read the reasoning. The result can be exported as text or CSV.

> **The single most important rule.** The measured area must contain the finding *and* the surrounding skin. Every feature is relative: if the ring lies entirely inside the finding there is nothing to compare against, and the app says so — "*area and surroundings share the same colour*".

3. How to photograph

The largest error is not a missing calibration card but the **phone's automatics**. White balance actively tries to remove the colour cast — precisely the discolouration you want to measure. A livid bruise can "heal" while still in the viewfinder.

Do this	Why
Lock exposure and white balance by long-pressing on **healthy skin beside the finding**	The automatics stop "correcting" the discolouration
Shoot perpendicular to the surface	An oblique view changes lightness and saturation and shrinks the finding
Even, diffuse light; avoid flash	Both glare and hard shadow distort the measurement — the app reports either
Include the surrounding skin in the frame	The ring needs a reference
Place a scale if you need dimensions	Without one the app cannot determine size
Take a series from several angles	Shadow on a curved surface can mimic a haemorrhage

4. Measuring

The measuring area

A tap sets the centre. The slider changes the radius. The app continuously advises whether the lesion is correctly framed, occupies too small a part of the area, or whether the area lies entirely inside the lesion.

Snapping to the lesion

Optional snapping moves the measured area to the colour centroid of the lesion — useful when your tap was slightly off. It never runs while the slider is being dragged: otherwise the loupe would jump around the image.

The loupe

The circular inset in the top-right corner shows the measured area magnified. It is purely informative and takes no touches; the magnifier button below it (from 3.0) doubles its size when you need to see detail — for instance pinpoint haemorrhages within an erythema.

Lesion dimensions

The record needs dimensions, not a share of the image. The app derives them from the scale (two points on a ruler, automatically from a T2/T3 target, or from 3.0 from **LiDAR** when photographing from the Camera tab — see below) only for a round finding; for an elongated one (wound, scratch, band) a diameter says nothing. Length and width can therefore be **entered manually** with *enter dimensions…* in the result panel — measured with a ruler or a forensic scale at the finding. The unit (mm / cm / inches) is chosen in Info → Length unit; internally everything is computed in millimetres. Manual dimensions apply to one photograph and go into the text record and the CSV.

Scale from LiDAR (from 3.0, iPhone Pro)

iPhone Pro (12 Pro and later) has LiDAR: it knows the distance of every point in the frame. Together with the focal length of the lens (the camera's calibration data) this gives the scale directly — *millimetres per pixel*

= *distance* ÷ *focal length* — with no tapping and no object in the frame. In the Camera tab the measured point shows *LiDAR 28 cm · area ⌀ 9.6 mm* continuously. The **Take photo (LiDAR)** button captures a real photograph (not a video frame), saves it to Photos with the depth data embedded and opens it in the Photograph tab with the scale already set — the lesion-diameter estimate from the outline is available at once and the source of the scale (*LiDAR, distance 28 cm*) goes into the record.

Limits: the scale holds for a plane perpendicular to the lens axis — on a curved limb photographed from the side, a dimension along the curvature is underestimated. With digital zoom no scale is provided from the photograph (the effective focal length changes, and a silently wrong scale is worse than none). Below 5 cm LiDAR does not measure. For an evidentially important dimension a ruler in the frame remains the reference; LiDAR gives a dimension where there would otherwise be none. This is not image recognition — LiDAR measures time of flight and the calibration is a lens table; the app still links neither Vision, CoreML nor ARKit.

Measured values

The bottom panel shows the colour name and the conclusion; the numeric values (HEX, RGB, L*a*b*, uniformity, margin, roughness and the deviations against the surrounding skin) unfold with the *Detail* button. The setting is remembered.

5. Calibration

Calibration lives under the half-circle icon, top left, and comes in three levels:

Tool	What it corrects
Grey card	The colour cast of the lighting. Simple and quick.
Colour chart (24 patches)	Additionally, the distortion with which the camera renders colours themselves. Tap the four corners of the patch grid in the order top-left → top-right → bottom-right → bottom-left; the app computes a homography and a 3×3 correction matrix in linear RGB.
T2 / T3 target (automatic)	The calibration target from the aPrehliadky project is found in the frame by itself — no tapping. **T3** (A4, circular fiducials): with a ColorChecker Classic Mini inserted in the cut-out, a full correction matrix is computed from its 24 patches; without it, white is balanced from the grey row. **T2** (QR fiducials): white balance from the grey row. The fiducials found and the patches read are drawn over the photograph so you can see what the app read.
Scale in mm	Pixels to millimetres — two points on a ruler placed in the frame.

> **Calibration always applies to a single photograph.** Carried over to another image it will ruin the colour. It is therefore cleared when a new photograph is loaded, and the app says so.

The T2/T3 target is located without image recognition. The T3 circles are a threshold and connected components; the T2 QR codes are *not decoded* — they are found as dark textured squares, and their assignment is decided by the grey row, which only the correct geometry hits. Even for this feature the app uses neither Vision nor CoreML. The target is printed at 100 % on A4; print scale has no effect on colour, but it does affect the mm scale — so verify the print with a ruler (T3: circle centres A–B = 247 mm).

Print the target straight from the app: Info → Calibration → *Target T3 for printing (A4 landscape, PDF)* or *Target T2 for printing (A4, PDF)*. The PDF opens in the share sheet — send it to a printer or save it. T3 has a cut-out for a ColorChecker Classic Mini; T2 works with an ordinary printer (on home prints the grey row is more reliable than the colour patches, which are indicative only).

Why the chart matters for bruises in particular: the only time indication supported by the literature is the shift towards yellow — and the yellow-blue axis is also the axis with the largest rendering error. Without calibration the app is therefore more cautious about dating.

6. What the app measures

Twelve values are read from every area. They are the same numbers that make up the reference fingerprints (chapter 13), so measurement and reference are compared in one space.

Feature	What it expresses
Δ area lightness (median)	how much darker the whole area is than the surrounding skin
Δ core lightness	the same for the darkest component
core hue · core hue shift	what colour the core is and how far it is from the hue of the skin
Δa (red–green)	shift towards red
Δb (yellow–blue)	shift towards yellow or towards blue
colour deviation	total distance from the surrounding skin
pale rim rise · pale rim chroma gain	what the pale component does at the edge of the finding
surface roughness	texture — trustworthy only at sufficient resolution
edge contrast	sharpness of the margin
scale area	the proportion of pale scaly points

7. Method: the chromophore decides

The app does not conclude from the absolute colour but from **which substance causes it**. This is the core of the whole method and the reason it can separate findings that look alike.

Chromophore	How it shows
Melanin	Darkens the whole area, the hue stays yellow-brown (core 38–88°), the skin surface is unchanged. **It adds the yellow component.**
Haemoglobin	Adds colour but hardly darkens the skin — the blood is *beneath* intact skin. A livid haemorrhage has b* **lower** than the surrounding skin; melanin can never do that. **It removes the yellow component.**
Haematin (dried blood)	A very dark core, red-brown hue below 36°.
Missing pigment	Lighter and less saturated than the surroundings.
Fibrous tissue (scar)	Recognised by what it *lacks*: no dark component, a smooth surface without skin markings.
Exogenous pigment	A hue of 160–280° at a chroma of C ≥ 14 cannot be produced by any skin chromophore — the pigment must have been introduced from outside.

> **Melanin adds yellow, blood takes it away.** One sentence that decides most often in this app. Solar lentigo has Δb* +6.7 and seborrheic keratosis +8.4; haemorrhages −3.5 to −12. That is why a pigmented lesion can be separated from a bruise even when their lightness is almost identical.

8. Measured reference values

The thresholds are not estimated. They are computed on verified, openly licensed clinical photographs (Wikimedia Commons, August 2026) by a replica of the very measuring procedure the app uses. Rows marked *(photo)* come from a verification series of ten photographs and are locked by regression tests.

Finding	ΔL* core	core hue	ΔL* median
crust	−38 … −45	17–31°	−3 … −13
naevus	−13 … −23	54–63°	−1 … −7
healing abrasion	−10	26–33°	−1 … −3
scar (mature)	−2	–	+2 … +7
healthy skin	−4 … −7	52–54°	0 … −3
seborrheic keratosis	−25 … −44	53–64°	−11 … −16
verruca (hyperkeratosis)	+3 … −23	–	−3 … +19
livid bruise (photo)	−15 … −19	–	−7 … −10.3
resorbing haematoma (photo)	−23	65°	−11
petechiae (photo)	−14	40°	−1
Mongolian spot (photo)	−7	–	−3 (Δb* −3.3)

Finding	ΔL* core	core hue	ΔL* median
maturing scar (photo)	−5.5	–	+1.1
solar lentigo (photo)	−9.4	84°	−3.3 (Δb* +6.7)
psoriatic plaque (photo)	−25	29°	+3 (scales 9 %)
cherry angioma (photo)	−46	22°	−36 (Δa +37)

Seborrheic keratosis is separated from a naevus by the depth of the core (boundary −24). A verruca is carried by a pale, weakly saturated component (rim ΔL* +18 … +29 with a chroma loss of −7 … −15) with Δa below zero — keratin has no blood.

9. Assessment modes

Mode	What it does
Colour only	Measures and names the colour. No pattern assessment.
Forensic assessment	Injuries, their morphology, vitality, dating and differential considerations; postmortem changes when the subject is marked as deceased.
Dermatological assessment	Skin findings, the ABCDE criteria as published educational criteria, differential considerations.

> **Non-traumatic findings in forensic mode.** Naevus, lentigo, depigmentation, keratosis, verruca, Mongolian spot and haemangioma *must* be recognised — this protects against mistaking them for a bruise. Their percentage is nevertheless capped at 60 %, the heading is marked "non-traumatic finding", and the app refers you to the dermatological mode.

Postmortem lividity is offered **only** for an explicitly deceased subject (never for "not specified") and never in dermatological mode. The state of the subject is the one field the examiner always knows.

10. Clinical context and examination questions

A photograph does not carry the features that decide most often. The app therefore asks — and the answers can not only exclude a finding but also establish one.

Clinical context

- **State of the subject** — living / deceased / not specified
- **History of trauma** — yes / no / unknown
- **Duration of the finding** — days / weeks / months / years / unknown
- **Body region** (from 3.0) — head/face · neck · trunk · upper arm/forearm · wrist/hand · thigh/shin · knee/elbow · ankle/foot. A photograph of one patch of skin carries no information about where on the body it lies, and the app does not infer it from the image — you enter it. Location excludes whole

groups of findings more reliably than colour: a ligature or restraint furrow away from the neck and limbs is impossible (excluded outright on the head and trunk, weight 0.15 on a knee or elbow), abrasions are typical over joints, a Mongolian spot is lumbosacral, warts sit on hands and soles. The weights are deliberately mild — location *orders* findings, it does not create them.

From the examination (optional)

- skin in the area — intact / broken / crusted / healed
- blanches under pressure (diascopy) — yes / no / not tested
- skin markings in the area — preserved / absent
- level vs surrounding skin — level / depressed / raised
- bleeding from the lesion — yes / recently / no
- distribution — single site / dependent / patterned

The healing timeline

Duration is not a switch but a sequence of overlapping phases. A scar runs along that axis in the opposite direction to everything else — its probability *increases* with time, because it is a permanent sequela.

Factually impossible combinations are excluded outright: bruise, crust, wound or abrasion at a duration of months or years · wounds without trauma · a scar within a few days. An excluded finding is removed from the list *entirely* and the reason is stated — a merely soft weight could, after normalisation, give the single remaining option 100 %.

> **The clinical context is cleared when a new photograph is loaded.** Carried over to another finding it would silently distort the conclusions.

11. Percentages — what they mean and what they do not

A percentage expresses how well the measured numbers agree with reference descriptions from the literature. The score is rule-based and normalised.

- **It is not a probability of disease.**
- **It is not a diagnosis** and not a classifier's confidence.
- **Below 50 % nothing is "recognised".** If even the first option does not reach half, the headline reads *Uncertain finding — options only* and the two best are listed. Origin: a reddened elbow (erythema) came out as "band-like mark — a ligature furrow?" at 30 %, and the headline *Recognised pattern* turned that into a certainty the measurement never had.
- **The ceiling is 85 %** — no result may be read as certainty.
- The list narrows according to the context you enter; whatever is excluded is not listed.

Every candidate comes with reasoning: which measured features support it, what casts doubt on it, what should be verified by inspection or diascopy, and what it can be mistaken for. **The reasoning matters more than the number.**

12. Atlas of findings

The reference section: 67 findings with original schematic drawings, in both languages. Each card gives the appearance, colour, margins, texture, mechanism of injury, typical causes and a forensic note.

The categories cover mechanical skin injuries, thermal and chemical damage, postmortem changes, vascular and haemorrhagic findings, pigmented lesions, inflammatory and scaly dermatoses, benign skin growths and body modifications. Search works **in both languages at once** — an English name finds the Slovak card and vice versa.

> The schemas are simplified drawings of the typical appearance, not photographs of particular cases. This is deliberate: a photograph of one case tempts the reader to look for the finding by its resemblance to that one image.

13. Reference fingerprints

A fingerprint is a vector of twelve features (chapter 6) **measured on a verified clinical photograph** by a replica of the app's own measuring procedure. Eighteen are built in.

The comparison is a weighted distance in the feature space: every feature has its own tolerance and its own weight, both written out in the source code. The core hue is counted only when the core is distinctly dark — in a shallow core it is random. The result is a match score between 0 and 1.

How a match is applied

- It **confirms** a candidate that the rule branches have already offered.
- It **establishes** a candidate the branches overlooked — but only on a strong match and only through the same gates of mode and context.
- Clinical context and hard exclusions apply equally; a fingerprint **never overrides them**.

The healthy-skin guard

One fingerprint is measured on healthy skin. If a measurement resembles it more than it resembles any finding, no match is reported. Without this guard, weak fingerprints — a maturing scar, say, which is almost the colour of skin — would report a match on any skin at all.

Agreement of several examples

When several independent fingerprints of the same finding take part in a match, the score rises slightly (+3 % per example, at most +8 %) and the count is written into the reasoning. Three agreeing examples are a better-supported conclusion than one lucky hit.

14. Expert mode: learning from your own cases

An optional feature, **off by default**. It is switched on in Info and adds *Learn from this measurement…* to the assessment. It stores a measurement as a new reference fingerprint under your own label.

> **Twelve measured numbers and your note are stored — never a photograph and never data tied to an identifiable person.** The fingerprint joins recognition immediately and can be deleted individually.

From version 2.8 the fingerprint also records the **name of the person measuring** and whether the measurement was calibrated with a colour chart. Both travel with the file — without them, after a few rounds of exchange there is no way to tell whose record it is or what light it was made in.

> **Label only conclusions you have verified** — by autopsy, histology or course. A mislabelled fingerprint will pull the app in the wrong direction at every subsequent measurement. That is why the feature sits behind a switch.

Is this machine learning?

In the broad sense, yes: comparison against stored examples is a nearest-neighbour method ("lazy learning"), and that is a family of machine-learning methods. It is not, however, a **trained model**. A neural network learns by rewriting millions of internal weights that nobody can then read; here not one is rewritten — a row is simply added, one you can open, read, send to a colleague or delete. The app links neither CoreML nor Vision and there is no model in the bundle.

For forensic practice the difference matters: when the app says "contusion — 85 %", you know exactly why — these numbers, by this procedure, against this fingerprint from this photograph.

15. Sharing fingerprints with colleagues

Fingerprints travel as a `.ch3lab` file (JSON with a header, its own file type, sent by e-mail, AirDrop or through Files). An export carries **the whole drawer** — yours and those you previously received — and merging skips duplicates by record identifier. A file circulating among colleagues therefore accumulates and is never impoverished.

Packs

One import is one **pack**: a named set that can be disabled (it stays stored but takes no part in recognition) or deleted as a whole. If it turns out that a colleague had a whole series mislabelled, you revoke it with a single tap.

Quarantine

Received fingerprints **take no part in recognition** until you have gone through them. Without that, someone else's fingerprint would start changing your conclusions the moment the file was opened.

Review before acceptance

For every waiting fingerprint you can see what will change once it is accepted:

Verdict	Meaning
adds nothing	you already have an almost identical fingerprint under the same label

Verdict	Meaning
overlaps another finding	it lands on a stored fingerprint of a different label — the app will stop telling them apart at that point
will have no effect on recognition	it measures almost like healthy skin; the guard fingerprint suppresses it
confirms the current conclusion	the app already reports the same finding at that point
will shift recognition	the app currently reports something else there — which is exactly what learning is for

> **The app does not judge whether a label is correct.** That does not follow from the measured numbers — and it is honest to say why. The feature space is not clustered by label: two examples of *the same* finding (a fresh livid haemorrhage and a resorbing haematoma) can have a similarity of 0.003, while two *different* findings (a mature scar and healthy skin) reach 0.938. A check of the form "do you fit your own label?" would therefore discard honest examples and wave nonsense through. The app states only what will happen — the decision stays with you.

An honest caveat about a shared bank

Sharing multiplies errors too. Every fingerprint is measured on a photograph taken with an unknown camera under unknown light; when ten of you pool them, you pool ten different colour shifts as well. The agreement of several examples and the calibration flag dampen this but do not remove it. **Keep packs small and known** — a few colleagues whose photographic routine you know — rather than "the more the better".

16. Limits of the method

What the app cannot do	Why
Determine the age of a bruise	Colour tables assigning a number of days have not been confirmed in more recent work. The only supported indication is the yellow component: it is not observed in a haemorrhage younger than roughly 18 hours.
Separate jaundice from warm lighting without calibration	It is a whole-field discolouration — relative features cannot carry it.
Separate erythema from a fresh haemorrhage	They are indistinguishable by colour; diascopy decides.
Conclude a finding from a single photograph	Skin integrity, diascopy, skin markings and relief exclude whole groups of findings more reliably than any colour measurement.
Replace histology	Colour is a surface property. It says nothing about depth or structure.

Still unverified on a real finding: cyanosis, putrefaction, electrical mark, blister, pustule and molluscum. The app recognises them from physiology and the literature and says so openly. Colour-chart calibration is currently available only in the Photo tab; the Camera tab offers the grey card.

17. Privacy and data

- **No network access.** The app has no account, no backend, no analytics and no third-party library. Everything runs on the device.
- **No data collection.** A privacy manifest declares its only required-reason API (UserDefaults, category CA92.1, for the app's own settings), no tracking and no collected data types.
- **Photographs are neither stored nor transmitted.** The camera samples colour at the point you choose; the app records and shares no photo or video.
- **Learned fingerprints** live in the app's documents folder as readable JSON. They contain numbers and your notes — never an image. Keeping data tied to an identifiable person out of the note is the user's responsibility.
- **The `.ch3lab` file you send a colleague** contains those same numbers, your notes and the name of the person measuring. You can open it in any text editor before sending.

18. References

1. Saukko P, Knight B. Knight's Forensic Pathology. 4th ed. CRC Press, 2016.
2. DiMaio VJM, DiMaio D. Forensic Pathology. 2nd ed. CRC Press, 2001.
3. Payne-James J, Jones R (eds). Simpson's Forensic Medicine. 14th ed. CRC Press, 2019.
4. Langlois NEI, Gresham GA. The ageing of bruises: a review and study of the colour changes with time. Forensic Sci Int. 1991;50(2):227–238.
5. Pilling ML, Vanezis P, Perrett D, Johnston A. Visual assessment of the timing of bruising by forensic experts. J Forensic Leg Med. 2010;17(3):143–149.
6. Maguire S, Mann MK, Sibert J, Kemp A. Can you age bruises accurately in children? A systematic review. Arch Dis Child. 2005;90(2):187–189.
7. Bariciak ED, Plint AC, Gaboury I, Bennett S. Dating of bruises in children: an assessment of physician accuracy. Pediatrics. 2003;112(4):804–807.
8. Munang LA, Leonard PA, Mok JYQ. Lack of agreement on colour description between clinicians examining childhood bruising. J Clin Forensic Med. 2002;9(4):171–174.
9. Hughes VK, Ellis PS, Burt T, Langlois NEI. The practical application of reflectance spectrophotometry for the demonstration of haemoglobin and its degradation in bruises. J Clin Pathol. 2004;57(4):355–359.
10. Randeberg LL, Winnem AM, Langlois NEI, et al. Skin changes following minor trauma. Lasers Surg Med. 2007;39(5):403–413.
11. Bolognia JL, Schaffer JV, Cerroni L. Dermatology. 4th ed. Elsevier, 2018.
12. Abbasi NR, Shaw HM, Rigel DS, et al. Early diagnosis of cutaneous melanoma: revisiting the ABCD criteria. JAMA. 2004;292(22):2771–2776.
13. Grob JJ, Bonerandi JJ. The "ugly duckling" sign. Arch Dermatol. 1998;134(1):103–104.
14. Luo MR, Cui G, Rigg B. The development of the CIE 2000 colour-difference formula: CIEDE2000. Color Res Appl. 2001;26(5):340–350.

The reference criteria derive from the literature above. Older textbook tables assigning a specific number of days to a colour have not been confirmed in more recent work — the app therefore does not reproduce them and, instead of a date, states what is excluded.

Colorime3ka — professional manual · app version 3.0 · September 2026

The app is not a medical device. The decision always belongs to the physician.