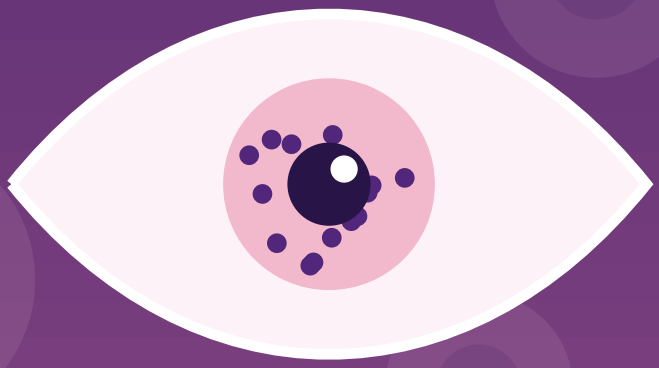

littlEye AI

User guide for physicians

version 3

Training and evaluating AI models from histology slides

iPhone · iPad · Mac — entirely on-device

Forensika

OBJASŇUJEME · SPÁJAME · POSÚVAME

forensika.eu

A project of three institutes, Comenius University · app 2.2.3 (11) · August 2026

Contents

What is new in version 3

- Cases — scanner slides grouped under one case with history, autopsy conclusion and toxicology; one tap gives a concise summary for Word
- Slide sharing (.le3set) — datasets travel between colleagues with diagnoses and reviews included
- Education — a “what next after H&E” guide, special stains, IHC, cellular elements and pigments
- Magnifications from 4× to 1000×, a stain field (H&E, PAS, trichrome, silver…) and zoom when viewing
- Bilingual SK/EN with a professional / plain-language switch
- Archive backup to an external drive, and five app icons to choose from
- Learning from corrections — fix a wrong result and save the image into your dataset in one tap; the next training round learns from it
- Field-of-view crop when classifying — mark a single field on an overview image and only it is evaluated (matching the training scale)
- Tissue check in case assessment — the organ-recognition model flags a possible slide mix-up
- Encrypted transfer files — .le3AI and .le3set open only in littlEye AI; automatic, no passwords
- 27 organs and tissues and an atlas of 90 diagnoses — the digestive tract, airways, endocrine glands, lymphoid tissue, muscle, vessels, bone and urogenital organs were added
- A colorectal cancer model on request — it distinguishes 9 tumour components; other models you train from your own slides (chapter 11)
- Model chaining when classifying — the app recognizes the tissue first, then runs only that organ's diagnoses

1 · What littlEye AI is

littlEye AI teaches artificial intelligence to look at histology slides through your eyes. You label microphotographs with diagnoses and the app trains a recognition model directly on your device, which then helps you assess new slides. Nothing leaves your iPhone, iPad or Mac: no server, no internet, no data collection.

What you will find in the app

- Dataset — H&E microphotographs of eight organs (brain, lungs, heart, liver, kidney, spleen, pancreas, skin)
- Slide review — colleagues assign diagnoses; every review is stored with a name (audit trail)
- Training — your own Core ML model with honest metrics (held-out test set, confusion matrix)
- Classify image — a preliminary diagnosis with a confidence figure and a draft report
- Cases — the autopsy workflow: slides + history + autopsy conclusion + toxicology → a concise summary
- Atlas and Education — descriptions of ~40 diagnoses, special stains, IHC, cellular elements, pigments
- Sharing — models (.le3AI) and slide datasets (.le3set) as a single file

Requirements

iPhone (iOS 17+), iPad (iPadOS 17+; on-device training works best on M-chip iPads), Mac (macOS 14+, Apple Silicon). The app works fully offline — no internet, no account, no server. It is bilingual (SK/EN) with professional and plain-language terminology — switch with the globe icon.

2 · First launch — in three minutes

1. Open the app, choose a language and enter your name (stored with reviews and models — audit trail).
2. On the Overview screen tap Load sample model — a bundled demo model with 10 diagnoses.
3. The app opens Classify image. Pick any image from Photos or Files — an ordinary photograph is fine for exploring the interface, since the demo model is trained on synthetic patterns. Look at the result: diagnoses ranked by confidence, plus a preliminary diagnosis with a microscopic description and suggested conclusion.

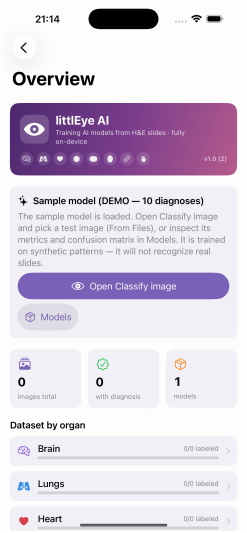

Overview — sample model, dataset of eight organs

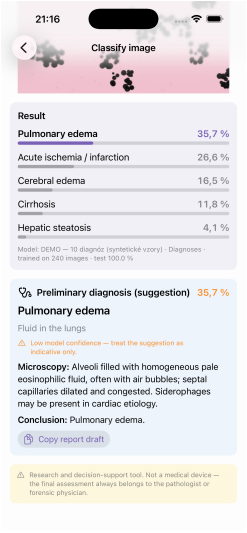

Classification: confidence, preliminary diagnosis, draft report

Important

The sample model is trained on synthetic patterns — it will not recognize real slides. It exists only so you can explore the interface. The app becomes genuinely useful once you train a model on your own images.

3 · Collecting images — the foundation

Your model will only ever be as good as your images. Six golden rules:

Six golden rules of imaging

- One magnification per model — we recommend 100× or 200×; the magnification is recorded on import and the app warns you if you mix them in training
- The finding should fill most of the field — the model judges the whole image
- Sharp, evenly exposed images with balanced white
- Avoid section edges, folds, wrinkles and staining artifacts
- Several images per slide, from different representative areas
- “Normal” is a class in its own right — the model needs it as much as pathology
- Every class from at least three different cases or blocks — otherwise the model learns one slide's hue instead of the tissue

A proven trap: the model learns the stain

In a trial model whose liver class came from a single block (deep purple), the model labelled redder liver from another block as kidney — because the kidney images happened to be red too. It had tied the hue to the organ. So collect images from several cases and ideally from different staining runs; accuracy grows with variety, not merely with counts.

Getting images into the app

- Microscope camera / slide scanner — export JPEG/TIFF and transfer via Files (USB-C drive, network share, iCloud) or AirDrop; the app needs no drivers
- iPhone with an eyepiece adapter — import photos straight from Photos
- In the dataset section pick the organ, choose the magnification in the toolbar and tap From Photos or From Files

4 · Labeling and reviewing with colleagues

1. Tap an image to open the detail and pick a diagnosis from the list (you can add your own). Pinch or double-tap to zoom into the image.
2. In Slide review colleagues go through images one by one and assign diagnoses. Every review is stored with the reviewer's name; the final diagnosis follows the majority.
3. Long-press a diagnosis to see its Atlas description — so everyone labels by the same criteria.
4. Marking a finding: if an image contains several findings, tap Mark area (crop) in the detail, drag over the region and save it as a separate training image. Cropping is how you tell the model WHERE to look.

Example 1 — a small lung model

Dr. Očko wants a model for three lung diagnoses. At 100× he collects 35 images of pulmonary edema, 32 of bronchopneumonia and 30 of normal lung — from twelve different autopsies. He and a colleague review each other's images in Slide review; four disputed slides are settled by a third vote. The dataset is ready for training.

5 · Training the model

1. Open Training, choose Organ diagnoses and the organ (or Tissue recognition for a model that tells organs apart).
2. Check the class list — the minimum is 2 classes of 5 images, we recommend 30+ per class. The app warns about class imbalance and mixed magnifications.
3. In Training parameters keep the defaults (25 iterations, ScenePrint v2, 15 % held-out test set).
4. Tap Start training. On an M1 iPad a few hundred images typically take 1–3 minutes; you can watch the progress and cancel it.

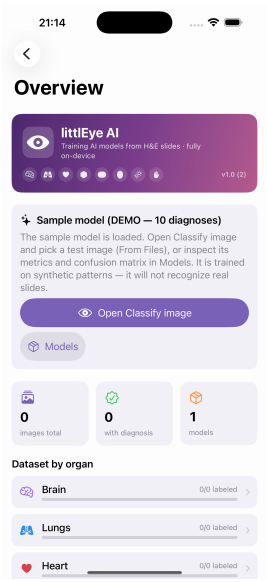

How to read the result

- Training / validation / test — the share of correctly classified images; test matters most (images the model never saw)
- Confusion matrix — row = true diagnosis, column = prediction; off-diagonal numbers show what the model confuses
- Sensitivity and precision for each class separately

A large gap between training and test accuracy means overfitting — add more varied images.

Example 2 — reading the confusion matrix

The model from Example 1 reaches 92 % on the test set. The confusion matrix shows it labelled 3 bronchopneumonia images as edema — both classes have congested septa. The fix: add bronchopneumonia images with prominent purulent exudate (crops from areas with neutrophils) and retrain. Afterwards the test accuracy rises to 97 %.

6 · Classifying an image in practice

1. Open Classify image, pick the model at the top and add an image (Photos or Files).
2. The result: diagnoses ranked by confidence. Confidence is always part of the result — never a bare claim.
3. For the top class the app shows a Preliminary diagnosis — a microscopic description and a suggested conclusion from the Atlas. Copy report draft puts the finished text on the clipboard (Word — Cmd/Ctrl+V).
4. Above the image you can see what the model was trained on (e.g. “100× · H&E”). If your image is a low-power overview, tap Select field of view (crop), mark an area the size of a single field and only it is evaluated — bringing structures to the scale the model knows. This alone fixes most nonsensical results.
5. If the model missed, pick the organ and the correct diagnosis right below the result in Learn from this image and save it into your dataset. The next training round learns from the correction —

this is exactly how the app grows on your material.

The rule

Below 80 % confidence the app itself warns that the suggestion is indicative only. AI is a second opinion, not a decision — the final assessment always belongs to the pathologist or forensic physician. And beware the opposite: high confidence on an absurd result means the image lies outside the model's world (it can only pick among classes it knows) — not that it is right. A detailed procedure for poor images is in the app: Methodology → “Getting the most accurate result from a poor image”.

7 · Cases — the autopsy workflow

The Cases section is for everyday work: all slides from one autopsy in one place, together with the case data.

1. Create a case with its protocol number (e.g. 680/2026).
2. Pick the organ and magnification in the toolbar and import the scanner slides; each can be viewed with zoom, assigned a diagnosis and a note.
3. Add the case history, the autopsy diagnostic conclusion and toxicology — briefly, key facts are enough.
4. Tap Assess case. The app evaluates every slide (your own diagnoses take precedence, unlabeled ones are judged by the model) and produces a concise summary.
5. Copy (for Word) — paste the summary into your report. Then archive the case.

Example 3 — a sample summary (shortened)

```
CASE 680/2026 — littlEye AI summary
SLIDE FINDINGS (n = 7)
• Lungs (3): pulmonary edema; congestion (model, 91 %).
• Heart (2): acute ischemia / infarction.
• Liver (2): hepatic steatosis (model, 88 %).
CONCLUSION — slides only
1. Acute ischemia / infarction. 2. Pulmonary edema, congestion. 3. Steatosis.
CONCLUSION — with case data
• The cardiac history supports the finding of acute ischemia.
• The alcohol history supports the finding of hepatic steatosis.
```

How the “conclusion with case data” works

The correlation is rule-based, not generative: the app looks for keywords in the data you entered (alcohol, hypertension, trauma, sepsis, heart failure, diabetes…) and pairs them with the slide findings. Findings without support in the data are listed separately. This is deliberately transparent and repeatable — easier to defend in court than text from a language model.

Tissue check — insurance against slide mix-ups

If a tissue-recognition model is imported (e.g. main_Sikuta.le3AI), the app cross-checks every case slide with it during assessment. When the model claims a different organ with ≥ 85 % confidence, the summary gains a WARNING — TISSUE CHECK section asking you to verify a possible mix-up. With scanner batches this is exactly the mistake that happens.

8 · Sharing: models and datasets

The app has two portable formats. The difference between them is the key to teamwork.

.le3AI — a finished model

Models & export → Export .le3AI. The file carries the model, its metrics, the confusion matrix and the authors' names — never the slide images. A colleague opens it with a double-click or via Import .le3AI and can use it immediately.

.le3set — a dataset of labelled slides

Dataset → sharing icon → Share slides (.le3set). It packages the currently displayed images (respecting the diagnosis filter) together with the diagnosis, magnification, colleagues' reviews and author. A colleague adds them via Import colleague's slides.

The files are encrypted — automatically

From version 2.0 both .le3AI and .le3set packages are written encrypted (AES-256). Only littlEye AI can open them; a modified or damaged file is rejected. Nothing to set up and no passwords — for a colleague with the app the file simply works, for anyone else it stays unreadable.

Why models cannot be “topped up”

Two terms that get confused: topping up would mean adding new images to a finished model — Core ML does not allow that; a finished model is a closed file that never changes. Retraining means building a new model from scratch from all images in the dataset — and that is exactly what the app does. That is why colleagues share images (.le3set), not models: imported images extend your dataset, and one tap on Start training produces a new, better model (minutes of work). The old model can be kept or deleted.

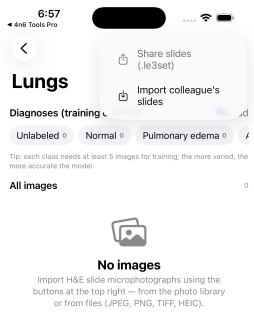

The sharing menu in the dataset — export and import in one tap

Example 4 — two physicians, one model

Dr. A has 40 images of pulmonary edema, Dr. B has 35 of bronchopneumonia. They swap .le3set packages over AirDrop. Each now has 75 labelled images, retrains, and test accuracy rises from 88 % to 94 %. Importing the same package twice reports “0 added, 40 were already in the dataset” — nothing is duplicated.

Data protection

A .le3set package contains only images and professional metadata — no case or patient data. Even so, share only microphotographs without identifying information; compliance with your institution's rules remains the user's responsibility.

9 · Diagnosis atlas and Education

Two knowledge sections that also make the app a textbook — usable without a single trained model.

Diagnosis atlas

About 40 diagnoses across eight organs. Each has a microscopic description, key features, a suggested conclusion (with a copy button) and, where relevant, a Next steps section — which special stains or IHC confirm or exclude the finding. The Atlas is also available while reviewing: just long-press a diagnosis.

Education — stains and IHC

- "What next after H&E" guide — 12 clinical situations: dating a hemorrhage, wound vitality, organization, the "silent" infarct, diffuse axonal injury, fat embolism, anaphylaxis, tumor of unknown origin, granulomas, sepsis, fibrosis and identifying brown pigment
- Special stains — Van Gieson, trichrome, reticulin, Perls, PAS, Sudan/Oil Red O, Congo red, Ziehl-Neelsen, Gram, Luxol, von Kossa, Giemsa
- IHC markers — CD68, fibronectin, CD62P, tryptase, C5b-9, troponin/myoglobin, β-APP, GFAP, CD31, cytokeratins, CK7/CK20 with organ markers, S100, CD45, vimentin, Ki-67
- Cellular elements — PMNs, eosinophil, basophil, mast cell, lymphocyte, plasma cell, monocyte/macrophage, siderophage, fibroblast, giant cells (with a timeline for dating)
- Pigments — lipofuscin, hemosiderin, melanin, formalin artifact, anthracosis, bilirubin

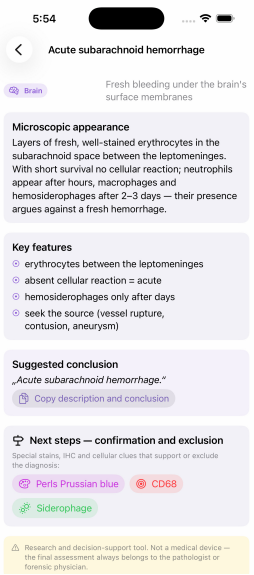

Atlas: description, features, conclusion and next steps

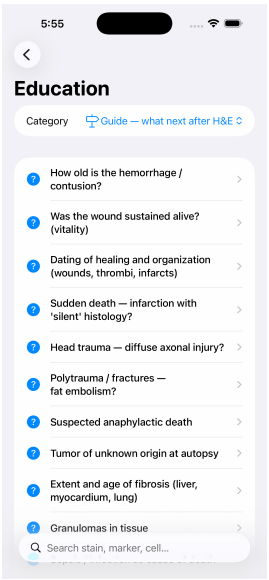

The guide — from question to method

The descriptions in the Atlas and Education draw on the referenced literature (Robbins & Cotran, Knight, Bancroft, Dabbs and others) — the full source list is in the app: About → References.

Example 5 — brown pigment in cells

You see brown pigment in the liver. Open Education → Guide → "Brown pigment in cells — what is it?" The app walks you through the decision: Perls blue = hemosiderin (old hemorrhage, ≥ 2–3 days). Perls negative with fine perinuclear granules = lipofuscin (age, cachexia). Birefringent across cells = formalin artifact. Tap the "Perls Prussian blue" chip to read the full method.

10 · Archiving and backup

Cases can be archived — they disappear from the everyday list but stay findable (Cases → … → Show archived).

Full archive backup: Cases → … → Back up archive to disk. Choose a folder (external drive, USB-C stick, network storage via Files) and the app writes a complete copy: the register, all images, all models and a manifest with the date and counts.

How the backup works

Every backup is a complete, self-contained copy of the entire archive as of that day (register, all images, all models) — not an increment since the last backup. Folders therefore grow with the dataset, but each one is complete on its own: any single backup suffices for a restore. Older backups can safely be deleted once the new one is verified.

A recommended routine

Back up to an external drive after every larger batch of images — at the end of the week, for instance. Backups are ordinary files, so they can be copied and archived with your institution's usual tools. The backup format is also groundwork for a future server: when one arrives, only the transport changes, not the data.

11 · Ready-made models and how to extend them

The app is a tool: it ships with a single sample model trained on synthetic patterns. Real models are ones you train on your own slides — and finished models can be exchanged as standalone .le3AI files, loaded in the app with the sparkle button in Models → Load your own model from a file.

A note on models from third-party collections

We have deleted the models we trained from the University of Michigan virtual collection. We asked their Tech Transfer Office for a position and the answer was that our use is commercial and the CC BY-NC-SA licence does not cover it, even where the images are not part of the app. We accepted that position. The same applies to models derived from TCGA-UT, under the same licence; we do not distribute those.

The colorectal model is trained from NCT-CRC-HE-100K (CC BY 4.0), which permits commercial use. It is unaffected by this question and is available on request at forensika.eu@icloud.com.

File	What it recognizes	Test
littlEye-hrubeCrevo-v1.le3AI	9 tissue components of colorectal cancer — adenocarcinoma, cancer-associated stroma, lymphocytes, mucus, smooth muscle, normal mucosa, necrosis, fat, background	88,3 %

How to extend a model with your own slides

1. Pick one organ you actually assess and create its classes — normal, plus the pathologies you want to distinguish. Fewer classes with enough images always beats many classes with a handful each.

2. Add your own microphotographs to the organs that interest you and assign diagnoses.
3. Run Training. A new model is created — that is your v3. Keep the original for comparison or delete it.
4. Export both the model and the dataset and send them to colleagues. Keep the version number in the name — otherwise, after a few months nobody can tell which generation they have.

What a model can and cannot do

A model recognizes only what it was trained on, and it must pick one of its classes — it cannot say "I have never seen this". Diagnoses come only from your own autopsy slides: 20–30 images per diagnosis, always from at least three different cases. Split training and test data by case, never within a case — otherwise the model memorizes one particular block and the accuracy figure lies. The confusion matrix in the model detail shows exactly which classes it mixes up.

12 · Vision: a shared "brain" for departments and universities

Every physician who labels slides in littlEye AI builds a piece of collective experience. The .le3AI and .le3set formats were designed from the start so those pieces can be joined — a personal tool can grow into a shared diagnostic memory for the whole field.

PHASE 1 — PHYSICIAN today	PHASE 2 — DEPARTMENT the next step	PHASE 3 — NETWORK the goal
Everyone trains locally on their own slides. Images travel as .le3set packages, finished models as .le3AI files — a colleague in Košice emails a model to a colleague in Bratislava.	The department appoints a curator. A collection of validated models grows on a shared drive — one "champion" per organ, validated on images from someone other than the person who trained it. Confusion matrices and the audit trail make quality auditable.	A Forensika server: a registry of models and anonymized datasets. Departments and faculties act as nodes — contributing annotations, downloading shared models, with disputed slides settled by consensus across institutes. The app's backup format is already prepared for that transfer.

The principles the shared brain rests on

- Consensus over authority — the final label comes from a majority vote; disagreements are the most valuable teaching material
- Cross-institutional validation — a model is measured on images from an institute that did not train it
- Audit trail — for every model it is traceable who trained it, when, on how many images and with what accuracy
- Anonymization — only microphotographs without identifying data are shared; models contain no images at all
- Teaching — students learn on slides with a consensus diagnosis; the Atlas and Education act as course notes inside the app

Why this matters

A pathologist's experience currently retires along with them. A labelled slide is experience that stays: it can be shared, measured, challenged and improved. A network of institutes exchanging models and annotations builds something no textbook can offer — a living, verifiable and continuously growing diagnostic memory for the field.

13 · Rules of use and limits

- littlEye AI is a research and decision-support tool — not a medical device; it does not provide diagnoses
- A model's output is always a second opinion with its confidence stated; the final assessment belongs to the pathologist or forensic physician
- A model trained at one magnification and stain may not work under different conditions
- Before practical use, validate the model on an independent set of images, ideally from a different reviewer
- Compliance of the images with institutional rules and patient data protection remains the user's responsibility
- The app has no network connection — data stays on the device, but you must make backups yourself

About the project

littlEye AI is developed in collaboration by specialists from three departments of the Faculty of Medicine, Comenius University in Bratislava: the Institute of Histology and Embryology, the Institute of Pathological Anatomy and the Institute of Forensic Medicine and Forensic Toxicology. The content is developed by specialists from these institutes.

The app and its ideas took shape over a long time — the content matured and was reworked based on experience from the apps 4n6 Tools and Colorime3ka (both on the App Store).

Contact and support

Ján Šikuta, MD, PhD., forensic pathologist · forensika.eu@icloud.com
littlEye AI — a Forensika forensic-medicine project · forensika.eu · version 2.2.3 (11)
Detailed guidance is inside the app: Trainer methodology, Hardware & capture, Atlas, Education.