From slide to model

The complete workflow: how to photograph a histology slide so that a usable model can be trained from it — and how to train, verify and export that model to colleagues in littlEye AI.

Why how you capture matters more than how many images you take

The model learns what you show it. If the images are blurred, at a different scale or with a different white balance, the model learns exactly those differences — not the tissue. Measurements on the littlEye AI baseline model showed that changing the magnification from 20× to 40× drops accuracy from 78 % to 52 %, and a poor phone-through-the-eyepiece image down to 30 %. An hour spent setting up the microscope therefore does more than a thousand additional bad images.

1 · What to capture with

The workflow works at all three levels of equipment. They differ in speed and comfort, not in the usability of the result.

Equipment	How it works	What to watch out for
Microscope + phone through the eyepiece	An adapter costing a few tens of euros holds the phone above the eyepiece. Enough to get started.	The greatest risk: the phone's automation changes exposure and colours between shots. Lock the exposure and the white balance. Do not use digital zoom.
Microscope + microscope camera	A camera in a C-mount adapter replaces one eyepiece; the image goes straight to the computer.	This is the sensible optimum. A 5–12 Mpx camera with a 1× adapter covers the field at the 20× objective with sufficient resolution.
Slide scanner	Scans the whole slide at once into one large file; patches are then selected on the computer.	The fastest and most consistent method — one slide yields hundreds of patches. The investment pays off only at dozens of slides per day.

An argument for the head of department: a microscope camera for roughly a thousand euros lets you collect data immediately. A slide scanner (tens of thousands) pays off once it is clear the material really is being collected and the models really are used. The sensible order is to start with a camera, show results, and only then ask for a scanner.

2 · The slide

The model cannot repair what the section ruined. A usable slide has:

- a thickness of 3–5 µm — a thicker section overlaps nuclei and loses detail
- even H&E; staining without overdone haematoxylin; nuclei must be blue-violet and readable, not black
- no folds, tears or bubbles in the assessed area
- a clean coverslip — dust and fingerprints show up in the image as artefacts the model will happily “learn”

If you collect material from several labs, expect the staining to differ. That is not a flaw — on the contrary, this very diversity makes the model robust. The condition is that every class has material from several sources, not that one class comes from one lab and another class from a different one.

3 · Microscope setup — Köhler illumination

It takes two minutes and is the single biggest quality improvement you can make for free. Incorrect illumination reduces contrast and resolution so much that even an expensive camera will not help.

1. Focus the slide with the 10× objective.
2. Close the field diaphragm (the lower diaphragm at the light source) almost completely — a bright polygon appears in the field of view.
3. Focus the polygon's edge by adjusting the condenser height until it is sharp.
4. Centre it with the two condenser screws into the middle of the field.
5. Open the field diaphragm just enough for the polygon to disappear beyond the field's edge — no more.
6. Set the aperture diaphragm (on the condenser) to roughly 70–80 % open. Closing it further raises contrast but lowers resolution and creates false contours.
7. Re-check points 2–5 whenever you change the objective.

> **Check**
>
> The background outside the tissue should be evenly white without a colour cast and without darker corners. If the corners are darker, the condenser is off-centre or the field diaphragm is too closed.

4 · Choosing the magnification

The table comes from measurements on the littlEye AI baseline model (19 tissues, held-out test on 286 images). The test images were modified to imitate other conditions:

Conditions	Accuracy	Note
20× objective (200× total), 480 px	78.3 %	the scale the model was trained at
40× objective (400× total)	52.1 %	structures are twice as large as the model expects
100× oil immersion (1000× total)	10.8 %	with 19 classes chance is 5.3 % — guessing
240 px resolution	69.9 %	still usable
160 px resolution	59.4 %	borderline
96 px resolution	29.7 %	unusable

Recommendation: the 20× objective (with a 10× eyepiece that is a total magnification of 200×) and a patch of at least 480 px on the shorter side. At this magnification the field shows both the tissue architecture and cellular detail — exactly the balance the model learns from best.

The 40× objective and oil immersion have their place for assessing nuclear detail by eye, but as training material for a model trained at 20× they do not work. If you want a 40× model, train it separately — mixing magnifications in one dataset is the surest way to ruin both.

5 · Scale — measure it once and be done forever

“20× objective” says nothing precise about how many micrometres one pixel covers — that depends on the camera, the adapter and the tube factor. It can, however, be measured once, and then it holds.

Procedure with a stage micrometer

1. Place a stage micrometer on the stage (a slide with an etched scale, usually 1 mm divided into 10 µm). If you do not have one, the histology institute almost certainly does.
2. Photograph it with every objective you will use — at the same camera and resolution settings you will use for the slides.
3. In any image viewer, read off how many pixels 100 µm occupies.
4. Divide: µm per pixel = 100 ÷ number of pixels. Write the number down next to the objective.

If 100 µm occupies…	µm per pixel	Field of a 480 px patch	Assessment
200 px	0.50	240 µm	ideal — matches the littlEye AI baseline model
400 px	0.25	120 µm	more detailed; downscale a 960 px patch to 480 px
100 px	1.00	480 µm	coarser; usable, but detail is missing
50 px	2.00	960 µm	an overview image — unsuitable for training

A scale bar directly in the image

Once you know the µm per pixel, you can put a scale bar — a line of known length — into the image. Most microscope-camera software does this automatically; if not, it can be added by hand in any editor.

> **Important: put scale bars only into images for publication and teaching**
>
> Do not put scale bars into training images. The model would learn to recognize precisely that line and the text next to it — and since the bar differs at every magnification, it would learn to guess the magnification from it instead of the tissue. It is the same trap as labels, arrows and circles in the image.
>
> The procedure that works: photograph the field once, save a clean file for training and a copy with the scale bar and caption for the lecture.

6 · Capturing

- Resolution: save at least 480 px on the shorter side of the assessed field; larger is better — you can always downscale, you can never add detail back.
- Format: TIFF or PNG if you can. JPEG is fine at 90 % quality or higher — but save it only once. Repeated JPEG saving adds artefacts that accumulate.
- Set the white balance once on an empty area of the slide and lock it. Automation changes colours between shots and the model interprets that as a tissue difference.
- Lock the exposure as well. The background should be bright but not blown out — the histogram must not have a spike at the very right edge.
- Focus on the nuclei, not on the coverslip or on dust.

- Switch off sharpening and noise reduction in the camera if possible. Both alter the texture the model learns from.
- Do not use digital zoom — it enlarges pixels, not detail.

7 · How many images, and from where

This is the part where the mistake is most often made — and shows up only at the end.

Rule	Why
At least 3 different cases per class, ideally 5 or more	Without this the model learns the colour and hue of one particular block. It genuinely happened to us: a liver trained from a single case classified a redder liver from another lab as kidney.
At least 100 patches per class, better 300	Below a hundred images the metrics are random — the model can look good by luck alone.
Different places on the slide, not ten shots of one field	Ten images of the same spot are, to the model, almost one image.
Balanced classes — a similar number of images for each	If one class has ten times as many, the model learns that guessing that class pays off.
One patient's images belong either to training or to the test — never to both	Otherwise the model memorizes the individual features of that tissue and the test shows an undeservedly high number. littlEye AI splits randomly, so add all images of one case at once, and for serious evaluation set a test case aside.

8 · Editing the image — what is allowed and what is not

Yes: crop to the assessed field (a square), rotate by a multiple of 90°, downscale to 480 px, convert formats without repeated compression.

No: filters and “enhancements”, extra sharpening, increased contrast or saturation, noise removal, labels, arrows, circles, scale bars, watermarks, collages of several fields in one file.

The reason is always the same: the model does not distinguish between the tissue and whatever you added to the image. If all tumour images are sharpened and all normal ones are not, the model learns to recognize sharpening.

9 · Import and labelling in littlEye AI

1. Open the Overview and choose an organ, or add one.
2. Import the files with the Add images button. Select all images of one case at once — they are easier to label.
3. Assign each image a diagnosis from the catalogue. The catalogue has over a hundred entries with microscopic descriptions and key features; if it lacks your diagnosis, add your own.
4. Fill in the magnification and stain. The model carries these along and warns you during evaluation when an image does not match.
5. Write the case or block number into the note. Later this tells you how many different cases a class really contains.

If you have a dataset from a colleague, import it on the Overview with the Import dataset (.le3set) button. Images you already have are not duplicated — each carries its own stable ID.

10 · Training

In the Training section choose the model type and start. The settings mean this:

Setting	What it does	Recommendation
Model type	Tissue recognizes the organ; diagnostic distinguishes diagnoses within one organ.	Start with a diagnostic model for the organ you actually assess.
Iterations	How many times the model passes through the data.	25 is a sensible baseline. More makes sense only with a large dataset — and watch whether only the training accuracy grows.
Augmentations	Artificially adds rotated and flipped copies of the images.	Enable rotation and flipping with a small dataset. Above a thousand images per class switch them off — real diversity is better.
Held-out test	Sets 15 % of the images aside; the model never sees them during training.	Never switch it off. Without it you have no way of knowing whether the model is any good.

11 · Reading the results

After training, three numbers and a confusion matrix appear. Only the third number matters.

- Training accuracy — how well the model handled what it saw. It is always high; it says nothing.
- Validation accuracy — a running check during training.
- Test accuracy — on images the model never saw. This is the only honest number.

How to read the confusion matrix

A row is the true diagnosis, a column is the model's answer. Numbers off the diagonal are confusions. What matters is not how many there are — it is which classes get confused and with what. A confusion that makes histological sense (tumour stroma for smooth muscle) is fixed by adding images. A confusion that makes no sense means something in the dataset is mislabelled — and that is a finding, not a malfunction.

Accuracy above 95 % on mixed material is a reason for suspicion, not for joy. Almost always one of the traps from chapter 7 is behind it.

In the Classify image section, try the model immediately on an image that was not in the dataset. If the result does not match your assessment, use Learn from this image: assign the correct diagnosis, the image is saved into the dataset and the next trained model will know it. This is the learning loop — the model is not fine-tuned, but the dataset grows and every next model is better.

12 · What it looks like in reality — the colon model

So that the numbers do not stay abstract: this is how a model trained in littlEye AI from the public NCT-CRC-HE-100K dataset (Kather et al., licence CC BY 4.0) turned out. Nine tissue components of colorectal carcinoma, H&E;, 20× objective.

Parameter	Value
training images	6,300 (700 for each of 9 classes)
test images	7,180 — from 50 patients the model never saw
settings	25 iterations, rotation + flip augmentation
training accuracy	94.9 %
test accuracy	88.3 %
training time	under a minute on an Apple Silicon MacBook

Per-class results

Class	Accuracy	Note
Background — no tissue	98.6 %	
Mucus	95.5 %	
Debris and necrosis	94.4 %	
Lymphocytic infiltrate	93.2 %	
Adipose tissue	91.4 %	
Normal mucosa	91.2 %	
Smooth muscle	83.3 %	
Adenocarcinoma	83.0 %	7.5 % confused with normal mucosa
Cancer-associated stroma	44.9 %	46.6 % confused with smooth muscle

What the single weak class tells us

Cancer-associated stroma was confused with smooth muscle in almost half the cases. That is not a malfunction — it is a histologically honest confusion. Both tissues are spindle-celled and eosinophilic, and in a field 112 µm wide not even a pathologist can reliably tell them apart without surrounding context. What decides is the arrangement into fascicles and the relationship to neighbouring glands — and that does not fit into such a small patch.

This is exactly how a confusion matrix should be read. A confusion that makes histological sense speaks about the limits of the method. A confusion that makes no sense speaks about a labelling error in the dataset. The first is fixed with a wider field of view, the second by fixing the data.

For comparison: the same app also trained a tissue model that recognizes 19 different tissues with a test accuracy of 80.6 %. Both models run on the same device, without an internet connection, and train in minutes.

13 · Export and sharing

1. In Models & export find the model and press Export .le3AI.
2. Name the file so that it is clear what it contains and which version it is — for instance main.le3AI for your principal model, or littlEye-srdce-v2.le3AI for a specific organ.
3. Record the provenance — the source of the images, the licence and the date. Without it the model cannot honestly be passed on a year later.
4. If you also share the images, export the dataset (.le3set). A colleague imports it, adds their own slides and retrains — that is how v2 and v3 arise and the field's shared knowledge grows.

Both files are encrypted and can be opened only by littlEye AI. The model carries its whole audit trail: author, date, image counts, metrics, confusion matrix, stain and magnification. Whoever receives it sees what it stands on.

14 · Checklist

Before you start training:

- Köhler illumination set, background evenly white
- 20× objective, µm per pixel measured and written down
- white balance and exposure locked
- patches at least 480 px, square, without scale bars or labels
- at least 3 different cases per class
- at least 100 patches per class, classes roughly balanced
- images from different places on the slide, not from one field
- every image has a diagnosis, magnification, stain and case number
- held-out test switched on

After training:

- test accuracy is below 95 % (higher is suspicious)
- confusion matrix reviewed, confusions make histological sense
- model tried on an image that was not in the dataset
- image provenance and licence recorded

Ján Šikuta, MD, PhD. · forensika.eu@icloud.com · littlEye AI 2.2.3 · August 2026