



# Imaging Workbench

## User Guide

A scan arrives as a DICOM folder or a NIfTI file, and the tools that can open it, segment it and measure it are scattered across a viewer, a scripting environment and a model repository that each need a specialist. The Imaging Workbench brings them into one browser: view and measure the scan, preprocess and register it, segment it by hand or with one of eleven verified AI models, turn a mask into a named biomarker, follow it across visits, and run the same recipe over a whole cohort, with every action recorded and a review decision on every result.

# Contents

---

|                                        |           |
|----------------------------------------|-----------|
| <b>Introduction</b>                    | <b>3</b>  |
| What it is                             | 3         |
| What it does                           | 3         |
| The problems it solves                 | 4         |
| System requirements                    | 4         |
| Before you start                       | 5         |
| The data in this guide                 | 5         |
| <b>The workspace</b>                   | <b>7</b>  |
| Before anything is imported            | 7         |
| Once a series is loaded                | 9         |
| <b>Importing a scan</b>                | <b>11</b> |
| The Import dialog                      | 11        |
| <b>Reading the volume</b>              | <b>13</b> |
| Navigating                             | 13        |
| Window, level, measurement and zoom    | 15        |
| Automatic quality control              | 16        |
| <b>Preprocessing and registration</b>  | <b>18</b> |
| The preprocessing pipeline             | 18        |
| Registering one series onto another    | 21        |
| <b>Segmentation and quantification</b> | <b>24</b> |
| The five ROI tools                     | 24        |
| Reviewing a segmentation               | 27        |
| Radiomics and biomarkers               | 27        |
| History                                | 29        |
| <b>AI models</b>                       | <b>31</b> |
| The model catalogue                    | 31        |
| Running a segmentation model           | 32        |

|                                                           |           |
|-----------------------------------------------------------|-----------|
| Benchmarking a model .....                                | 35        |
| <b>Detections and classification .....</b>                | <b>36</b> |
| Detection counts and the class filter .....               | 36        |
| Reviewing detections .....                                | 37        |
| Classification and its review .....                       | 38        |
| <b>Two more modalities .....</b>                          | <b>40</b> |
| X-ray classification .....                                | 40        |
| Parasite detection on a blood smear .....                 | 41        |
| <b>Longitudinal tracking .....</b>                        | <b>43</b> |
| Exporting a trajectory .....                              | 44        |
| <b>Recipes and batch .....</b>                            | <b>46</b> |
| Cohort export .....                                       | 48        |
| <b>A second anatomy model .....</b>                       | <b>50</b> |
| <b>DICOM and de-identification .....</b>                  | <b>52</b> |
| <b>Managing studies .....</b>                             | <b>54</b> |
| <b>Appendix A: Control reference .....</b>                | <b>55</b> |
| <b>Appendix B: Supported formats and modalities .....</b> | <b>57</b> |
| <b>Appendix C: The model catalogue .....</b>              | <b>58</b> |
| <b>Appendix D: Known limitations .....</b>                | <b>60</b> |
| <b>Appendix E: Data used in this guide .....</b>          | <b>61</b> |
| <b>Index .....</b>                                        | <b>62</b> |

# Introduction

---

## What it is

The Imaging Workbench takes a scan in DICOM, NIfTI or a plain research image such as a microscopy field or an X-ray, and gives you somewhere to work on it in a browser: view it slice by slice, measure it, preprocess and register it, segment it by hand or with an AI model, extract quantitative features, turn them into named biomarkers, track a participant over repeated visits, and export what you found. Eleven AI models for imaging work are admitted as of this guide. Each was converted to one verifiable runtime and each is described with its own measured accuracy.

It runs on your own machine and makes no outbound call of any kind. It is one of the four Novasyne analysis workbenches. Its biomarkers export as a tidy table that the Analytics Discovery Workbench reads directly, and as a biomarker template that a STUDYPRO study imports.

## What it does

- **Reads the scan as it is.** DICOM with its own geometry and modality tag, NIfTI volumes, and PNG, JPEG or TIFF images at their full bit depth, with a quality check on every import that names what it found.
- **Shows the volume.** Axial, coronal and sagittal planes, a slice slider, window and level, zoom, a ruler in real millimetres, and a compare pane for a second series beside the first.
- **Preprocesses and registers as new series.** Reorientation, resampling, bias field correction, denoising, normalisation and masking, and rigid or affine registration of one volume onto another, each writing a new derived series and never touching the original.
- **Segments five ways by hand and eleven ways by model.** Threshold, region growing, rectangle, polygon and brush, and a catalogue of admitted models for segmentation, detection and classification, every result landing in the same reviewable list.
- **Measures what a mask contains.** Volume, intensity and shape features from any segmentation, saved as named, versioned biomarkers with their own review status.
- **Follows a participant.** Repeated visits grouped by participant, a measured change in volume between them, and an export of the trajectory.
- **Runs the same recipe over a cohort.** Preprocessing, models, radiomics and biomarkers as one recipe replayed over any set of studies as a background job, with a cohort table at the end.
- **Records everything.** A History tab per series with every step and its parameters, and a review decision on every segmentation, detection, classification and biomarker.

## The problems it solves

**Opening the scan needs a specialist toolchain.** A DICOM series from the scanner, a NIfTI volume from a pipeline and a JPEG from a microscope each want a different viewer, and the segmentation model wants a Python environment somebody maintains. The workbench opens all three formats in a browser and runs every admitted model from one dialog.

**A model's result is a picture nobody checked.** A mask is drawn over the image, a count is quoted, a class is named, and nothing records whether a person agreed. Here every segmentation, detection, classification and biomarker carries a review decision, written into the study's provenance beside the result it judges, so the reader of an export can see which masks were accepted, which were corrected by hand and which were turned down.

**A measurement cannot be traced back.** A volume in a spreadsheet with no record of which mask it came from, which preprocessing ran first, or who accepted it. Every series here has a History with each step and its parameters, and a biomarker names the feature set, the mask and the definition it was computed from.

**The right data under the wrong label.** A NIfTI file whose axes are stored in an unusual order shows a sagittal profile under the Axial tab and nothing warns you. The workbench's quality check names the axis order it found and the step that fixes it.

**Fifty participants means fifty repetitions.** Load, preprocess, run the model, extract, save, fifty times. A recipe records what to ask for and the batch replays it over every study you tick, skipping by name whatever a step cannot take and writing the cohort table another workbench reads.

## System requirements

|                   |                                                                                                                                                                                     |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Processor</b>  | Intel Core i5 or higher, or the equivalent. Segmentation models run on the CPU; HD-BET takes about two minutes on a whole-head volume.                                              |
| <b>Memory</b>     | 16 GB of RAM or more. A volumetric series is held in memory in full, and the sliding-window models hold their patch stack beside it.                                                |
| <b>Disk space</b> | About 6 GB for the application, its environment and the model weights. Imported studies and their derived series are kept beside it, so allow several times the size of your scans. |
| <b>Browser</b>    | A current version of Chrome, Edge or Firefox                                                                                                                                        |
| <b>Network</b>    | None. The workbench makes no outbound call of any kind.                                                                                                                             |

## Before you start

|                        |                                                                                                                                      |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| <b>The application</b> | Running and reachable in a browser.                                                                                                  |
| <b>Your data</b>       | A scan in DICOM, NIfTI ( <code>.nii</code> , <code>.nii.gz</code> ) or a plain image (PNG, JPEG, TIFF) for X-ray or microscopy work. |
| <b>Nothing else</b>    | No API key, no account, no network.                                                                                                  |

Two things are true of this workbench and worth knowing before the first page.

**Nothing leaves the machine.** There is no assistant here, no external model call, no telemetry. A scan is read, held in memory, worked on and written back out; that is the whole data path. The exports this guide covers, to the Analytics Workbench and to STUDYPRO, write files to this machine's own disk.

**One series is viewed at a time.** The workbench holds a single active viewer session, and opening a different series replaces it. The compare pane shows a second series beside it for looking, not for editing.

### ONE WORKSPACE, ONE PERSON

The loaded series belongs to the application, not to your browser tab. Two tabs, or two people on the same machine, are two windows onto the same viewer. This is a tool for one person at a time working on one scan, which is why it binds to the local machine and refuses to start on an address other machines can reach unless it is told that sharing the one workspace is what you want. Where the NeuroStudio Gateway is installed in front of it, each signed-in person gets a workspace of their own and the refusal does not apply.

This guide is task-oriented. Each chapter answers a question you might arrive with, such as *why does the Coronal tab not show a coronal slice?*, *what does the workbench verify before it trusts a model?* or *how do I know if a detection is real?*, and shows the screens you will see. The appendices are the control reference, the supported formats, the model catalogue, the known limitations and the data used, and the Index at the back lists every term with the pages it appears on.

## The data in this guide

Every figure comes from real data, none of it synthetic and none of it patient data.


| Data                                | Used for                                                                                              |
|-------------------------------------|-------------------------------------------------------------------------------------------------------|
| NFBS repository, one T1 brain MRI   | Import, navigation, quality control, preprocessing, the five ROI tools, radiomics, biomarkers, HD-BET |
| NFBS, two visits of one participant | Registration, the longitudinal chapter, and two of the three studies in the batch                     |
| Prostate158, one T2 prostate MRI    | A second anatomy model with two structures                                                            |
| BCCD, one blood smear               | Detection, with the per-class count and filter                                                        |
| PatchCamelyon, one lymph node patch | Classification in true colour                                                                         |
| A chest X-ray and a wrist X-ray     | Two more classifiers on one modality                                                                  |
| An NIH malaria smear                | Parasite detection                                                                                    |
| A pydicom CT test file              | DICOM import and de-identification                                                                    |

Appendix E names each dataset.

# The workspace


---

## Before anything is imported



**Figure 1.** The workbench before anything is imported.


Import and Studies are the only controls that mean anything with nothing loaded. The rail down the left side is the whole application.



**Figure 2.**

Seven icons  
are the whole  
application.


Top to bottom: **Import** a scan, browse **Studies**, **Preprocess & Register** a series, run an **AI Model**, track a participant **Longitudinally**, run a **Recipe & Batch** over many studies, and **About**.



**Figure 3.** About names the build you are running.


**About** is the first thing to quote when you report a problem: the version, the build and the runtime the models are executed by.

## Once a series is loaded



**Figure 4.** A structural MRI, ready to work with.

Three regions. The **viewer** fills the left and centre. A **compare pane** on the right stays hidden until a second series is chosen to view beside this one. The **sidebar** is the whole working surface for everything this guide covers next.



**Figure 5.** Window and level, ROI tools, biomarkers, AI overlay, detections, classification and quality control all live in this rail.

Every section here reacts to what the loaded series actually has. Detections and Classification stay hidden until a model of that task has run, the AI overlay selector stays hidden until a segmentation exists, and Biomarkers says plainly that there are none yet rather than showing an empty table.

# Importing a scan

## The Import dialog

Import Imaging Data

PARTICIPANT REFERENCE (OPTIONAL)

DEMO-MRI-01

Free text only — this app does not connect to STUDYPRO directly.

MODALITY

MRI

Used when the source format does not state a modality (NIfTI, plain images). DICOM overrides this with its own Modality tag where recognised.

FILES

Drag and drop files here or click to browse

DICOM (.dcm), NIfTI (.nii, .nii.gz), images (.png/.jpg/.tif), or a .zip of any of these

sub-A00032875\_ses-NFB3\_T1w.nii.gz (9.4 MB)

Import Cancel

Figure 6. The Import dialog.

A **participant reference** is optional free text. This application does not connect to STUDYPRO or any other system directly, so nothing here is validated against a study roster. It is used to name export files and to group visits, so it may not contain a slash, backslash, quote, wildcard or colon, and it is capped at 64 characters; the dialog says so if you try.

**Modality** is used only when the source format does not declare one itself. DICOM's own Modality tag overrides it, but a NIfTI file or a plain image has to be told. It decides which models are offered for the series.

**De-identify** is checked by default. For DICOM it writes a separate de-identified derivative alongside the original rather than changing the file you uploaded; the last chapter shows it.

Import Imaging Data


Imported 1 series into study 544e292ebb95.

Import Cancel

Figure 7. A NIfTI volume brings in exactly one series.

A successful import shows how many series it found, then moves you to Studies. Anything in the upload that could not be used is named in a warning under that message and skipped, and the rest of the upload still comes in: a README inside a zip, a notes file, a DICOM that would not decode, a file whose name reduces to nothing safe, an archive that will not extract, or a NIfTI file with only

two dimensions, which is reported as one failed series while the rest of the upload is kept. Only an upload with nothing importable at all is refused. A 16-bit PNG or TIFF keeps its full intensity range on the way in; it is not squashed to 256 levels.



**Figure 8.** Every study imported so far, one card per participant.

Each card is one study: everything imported together, or everything sharing a participant reference through the longitudinal workflow. A study can carry many series, the original plus every derived segmentation, preprocessing result and registration this guide is about to produce. **View** on a series opens it in the viewer.

# Reading the volume

## Navigating

The toolbar's three tabs switch between the axial, coronal and sagittal planes of a volumetric series. The slice slider under the image runs through every slice on the current plane, the label beside it gives the slice index and the count, and the mouse wheel over the image steps through slices as well. Switching plane opens on the middle slice of the new plane.




Figure 9. Coronal.




Figure 10. Sagittal.

A plain image such as an X-ray or a microscopy field has one plane, and the plane tabs are hidden for it.





Figure 11. A slider through every slice on the current plane.

### THESE LABELS ASSUME THE FILE'S OWN AXIS ORDER MATCHES WHAT THE VIEWER EXPECTS

A NIfTI file's Axial, Coronal and Sagittal tabs are only guaranteed to show the plane they name when the file's own storage order already has one axis running left to right, one running posterior to anterior and one running inferior to superior. That is the RAS+ family of orientations, or any flip of it. A file whose axis order is **permuted** rather than merely flipped will show the right data under the wrong label, and nothing about the tab itself tells you that has happened.


This is not a hypothetical. The NFBS repository's own raw T1 files are exactly this case, and it is how this guide found the issue in the first place: the Axial tab opened on a sagittal profile. *Automatic quality control* below shows the warning the workbench gives for this, and *The preprocessing pipeline* shows the fix landing.

## Window, level, measurement and zoom



**Figure 12.** Width and centre narrowed to bring out soft-tissue contrast.

Width and centre are independent sliders. **Reset** returns to the range the series was loaded with and leaves whatever overlay you have switched on in place. The sliders take a step sized to the data, so a normalised volume whose whole range is 0 to 1 is as adjustable as a CT in Hounsfield units. Neither one writes anything. This is a display setting, not a preprocessing step.



**Figure 13.** A distance measurement, in real millimetres from the voxel spacing.

The ruler tool takes two clicks and reports the distance using the series' own voxel spacing, not screen pixels, so it is correct however far you have zoomed in.



**Figure 14.** Zoomed and panned to inspect one region closely.

## Automatic quality control

Every import runs a quality check automatically, without being asked.

### Quality Control

**Warning**

⚠ This file's axis order (PIR) is not in the orientation the viewer assumes, so the Axial/Coronal/Sagittal tabs will not show the plane they say they do. Apply "Reorient to canonical (RAS+)" in Preprocess & Register before relying on those labels.

- ✓ 192 slice(s), 256x256.
- ✓ All voxel values are finite.
- ✓ No blank slices.
- ✓ No duplicate slices.

**Figure 15.** A quality check runs automatically on every import.

This particular check is the one described above. The file's own axis order does not match what the plane tabs assume, so the check names the problem, names the actual axis order found, and names the specific fix: *Reorient to canonical (RAS+)* in Preprocess & Register, covered next. The

same four structural checks run on every file: slice count, finite values, blank slices and duplicate slices. A DICOM series gets one more. Its slices are ordered along the acquisition's own normal, so sagittal and coronal series stack correctly, and the spacing between slices is measured from their positions rather than trusted from a header; if the gaps are uneven, a **slice spacing** warning says by how much and which value was used.

#### **WARN, DO NOT REFUSE**

Quality control here flags problems; it does not block you from working with a file it has a concern about. A warning is information, not a veto. You decide whether it matters for what you are about to do.

# Preprocessing and registration

Preprocess & Register has two tabs, reached from the same sidebar icon. Both produce a **new** derived series; neither ever modifies the original you imported.

## The preprocessing pipeline

**Preprocess & Register**

Preprocess Register

INPUT SERIES

DEMO-MRI-01 — sub-A00032875\_ses-NFB3\_T1w.nii.gz (source, 192x256x256)


STEPS (APPLIED IN THIS ORDER)

- ☒ **Reorient to canonical (RAS+)**  
Reorients a volume to the closest canonical RAS+ orientation. No-op for plain images and for DICOM series imported before the workbench recorded their orientation.
- ☒ **Resample to isotropic spacing**  
Resamples the volume to a uniform voxel spacing using linear interpolation.  
target\_spacing\_mm: 1
- ☒ **N4 bias field correction**  
Corrects low-frequency intensity non-uniformity (MRI-oriented; SimpleITK N4BiasFieldCorrectionImageFilter).  
shrink\_factor: 4
- ☒ **Gaussian denoise**  
Gaussian smoothing filter.  
sigma: 1
- ☐ **Intensity normalisation**  
Z-score or min-max intensity normalisation.  
mode: zscore
- ☐ **Otsu foreground mask**  
Zeroes background voxels below an automatically-chosen Otsu threshold.


Run Cancel

**Figure 16.** Preprocessing steps, applied in the order listed.

Six steps are available, each independently switchable and independently parameterised: reorient to canonical RAS+, resample to isotropic spacing, N4 bias field correction, Gaussian denoise, intensity normalisation by z-score or min-max, and an Otsu foreground mask. Checked steps run in the order they are listed, not the order they were checked. A parameter outside a sensible range is refused before anything runs, with the step and the limit named: a resample spacing outside 0.1 to 10 mm, an N4 shrink factor below 1, or a negative denoise sigma.




**Figure 17.** Reorient to canonical, resample to 1 mm isotropic, then denoise.



**Figure 18.** A new derived series, not a modification of the original.


The result is a new series in the same study, visible from Studies or from the ROI/Segmentation panel like anything else derived from the original. Resampling records the spacing it actually achieved, not the one asked for: a whole number of voxels rarely spans the field of view exactly, so a 1 mm request can come back as 0.966 mm, and that achieved value is what the Details tab shows and what every later volume is computed from. The NFB3 scan is already 1 mm isotropic, so here

the spacing does not change. The success message names the new series and a short **pipeline version**, a fingerprint of exactly which steps and parameters ran, which the History tab records with it.




**Figure 19.** The same scan after Reorient to canonical: now a genuine axial slice under the Axial tab.

This is the direct fix for the axis-order warning above. The same NFBS scan, the same Axial tab, now showing a genuine axial cross-section, with orbits, sinuses and brainstem left-right symmetric, because the file's storage order has been brought into the orientation the viewer's plane labels assume.



**Figure 20.** The same checks on the reoriented series: a plain pass.

Quality control on the new series passes cleanly, because there is nothing left to warn about.



**Figure 21.** The reoriented series beside the original it was made from.


**Compare**, the split icon in the viewer toolbar, opens the compare pane and offers every series of every study. Here the original sits on the right as it was imported, showing its sagittal profile under an axial label, and the reoriented series sits on the left. Each pane has its own plane buttons and slice slider, so two visits, or a scan and its preprocessed copy, can be read against each other. The compare pane is for looking; every tool in the sidebar works on the left-hand series only.

### REORIENT FIRST, FOR ANY FILE THAT IS NOT ALREADY CANONICAL

If a source's own axis order is unfamiliar or unverified, applying *Reorient to canonical* before relying on the plane tabs costs nothing and removes the ambiguity entirely. A DICOM series is stacked from its own per-slice geometry at import, and the workbench records its orientation as it does for NIfTI, so the step works on DICOM too, as the step's own description in the dialog says. A DICOM series imported by an earlier version of the workbench has no orientation recorded; the step leaves it as it is, and importing it again is the fix. Plain images have no orientation, and the step does nothing to them.

## Registering one series onto another


The Register tab aligns a **moving** series onto a **fixed** reference, rigid or affine, and like preprocessing writes the result as a new series rather than altering either input.



The dialog box is titled "Preprocess & Register" and has two tabs: "Preprocess" and "Register". The "Register" tab is active. It contains three sections: "FIXED (REFERENCE) SERIES", "MOVING SERIES", and "TRANSFORM". Each section has a text input field. The "FIXED (REFERENCE) SERIES" field contains "DEMO-MRI-01 — sub-A00032875\_ses-NFB3\_T1w.nii.gz (source, 192x256x256)". The "MOVING SERIES" field contains "DEMO-MRI-01 — sub-A00032875\_ses-NFB3\_T1w.nii.gz (source, 192x256x256)". The "TRANSFORM" field contains "Rigid". At the bottom right, there are two buttons: "Run" and "Cancel".


**Figure 22.** Registering one volumetric series onto another.

Both dropdowns list every volumetric series across every study, which is what makes registering two *different* studies' scans possible from here. That is the real use, and the two visits of the longitudinal participant are the pair used in this guide.



The dialog box is titled "Preprocess & Register" and has two tabs: "Preprocess" and "Register". The "Register" tab is active. It contains three sections: "FIXED (REFERENCE) SERIES", "MOVING SERIES", and "TRANSFORM". Each section has a text input field. The "FIXED (REFERENCE) SERIES" field contains "A00061709 — sub-A00061709\_ses-NFB3\_T1w\_RAS.nii.gz (source, 256x256x192)". The "MOVING SERIES" field contains "A00061709 — sub-A00061709\_ses-NFB3\_T1w\_RAS.nii.gz (source, 256x256x192)". The "TRANSFORM" field contains "Rigid". At the bottom right, there are two buttons: "Run" and "Cancel".

**Figure 23.** The earlier visit as the fixed reference, the later visit as the one that moves.



The dialog box is titled "Preprocess & Register" and has two tabs: "Preprocess" and "Register". The "Register" tab is active. It contains three sections: "FIXED (REFERENCE) SERIES", "MOVING SERIES", and "TRANSFORM". Each section has a text input field. The "FIXED (REFERENCE) SERIES" field contains "DEMO-MRI-01 — sub-A00032875\_ses-NFB3\_T1w.nii.gz (source, 192x256x256)". The "MOVING SERIES" field contains "DEMO-MRI-01 — sub-A00032875\_ses-NFB3\_T1w.nii.gz (source, 192x256x256)". The "TRANSFORM" field contains "Rigid". Below the "TRANSFORM" field, there is a green status bar with the text "Registered series created (69940df0bc), final metric -0.2338." At the bottom right, there are two buttons: "Run" and "Cancel".


**Figure 24.** The registered series, and the final value of the similarity metric.

The reported metric is the optimiser's own final similarity score. Lower is a better rigid-body alignment for the metric this application uses. It is a number to compare across attempts on the *same* pair, not a universal pass threshold.

# Segmentation and quantification

## The five ROI tools


The ROI/Segmentation panel offers five ways to build a mask, all producing a reviewable segmentation series exactly like an AI model's output does. This chapter runs all five on the NFBS volume.



The screenshot shows the 'ROI / Segmentation' panel. It features a text input field labeled 'Threshold (whole volume)'. Below this is a row of four input fields: 'Low' with the value '0', 'High' with the value '1000', and two empty fields. A dark blue 'Run' button is positioned below these fields. At the bottom of the panel, it says 'No segmentations yet.'

**Figure 25.** Threshold segments the whole volume by intensity in one step.

**Threshold** takes a low and high intensity bound and segments the entire volume in one pass, the fastest of the five. Its mask is the one the radiomics and biomarker sections below are computed from.



The screenshot shows the 'TOLERANCE' tool interface. It features a blue circular seed point. Below the point, it says 'Click a point on the image to grow a region from it.'

**Figure 26.** One click grows a region from that point.


**Region growing** starts from a single clicked point and grows outward by intensity similarity, with a tolerance slider controlling how far it spreads. A seed that grows an empty region is refused with a message in the panel rather than saved as a mask of nothing.

**Rectangle** segments everything inside a dragged box on the current slice, the fastest way to isolate one clear structure. A click without a drag creates nothing; only a box with area becomes a segmentation.



**Figure 27.** Clicking out a boundary point by point.

**Polygon** builds a boundary from clicked points, finished explicitly with **Finish** rather than closing itself automatically, for an irregular region none of the other four tools would capture cleanly.



**Figure 28.** Freehand painting, with undo and redo.

**Brush** paints freehand at an adjustable size, with a full undo and redo stack, for a shape that genuinely needs a hand-drawn boundary. **Finish** sends the painting to the workbench. If it is refused, because nothing was painted, say, your strokes stay on the canvas with the reason beside them, and the canvas clears only once the segmentation has actually been saved.


The image displays a vertical stack of four segmentation review cards in the NovasynE interface. Each card represents a different segmentation method and includes the following elements:

- Method Name:** threshold, region\_growing, manual\_rectangle, and manual\_polygon.
- Status:** A dark blue badge in the top right corner of each card indicates the status is 'unreviewed'.
- Actions:** A row of buttons below the method name: 'View' (blue outline), 'Measure' (light blue), and a text input field containing the value '32'.
- Review Options:** A second row of buttons: 'Radiomics' (light blue), 'Accept' (light blue), 'Modify' (red text), and 'Reject' (red text).

**Figure 29.** Every segmentation made on this series so far, each reviewable on its own.


Every segmentation this series has produced, by any of the five tools or by an AI model, accumulates in one list. Each row offers to **View** it, **Measure** it, extract **Radiomics** from it, and record a review: **Accept**, **Modify** or **Reject**. The number beside Radiomics is the count of intensity bins the texture features are computed over, 32 by default; change it before extracting if your protocol fixes a different discretisation, because texture features are not comparable across bin counts.

## Reviewing a segmentation



**Figure 30.** Accepted: recorded in the study provenance, not just the screen.

Accepting or rejecting is not cosmetic. It is written to the study's provenance record and to the series' own review status, and it appears in the History beside the segmentation it judges. The workbench records the decision; it does not act on it. A rejected mask is still offered to a trajectory and still present in an export, with its status beside it, so that the person building the comparison decides what to leave out and the record shows what they decided.




**Figure 31.** Rejected: the brush segmentation is kept, marked, and left out of every trajectory and export.

A rejected segmentation is not deleted. It stays in the list and in the History with its decision, so that a reviewer later can see what was tried and turned down. **Modify** is the third decision, for a mask a person corrected before accepting it: the record says the result was changed by hand rather than taken as the tool produced it, which is the distinction a methods section has to draw.

## Radiomics and biomarkers


**Radiomics** turns a mask and its source image into a fixed set of shape, first-order intensity and texture features: a quantitative summary of the region, not a picture of it.



**Figure 32.** Shape, volume and intensity features, computed from the mask and the image.


The row shows three of the features and **Export CSV** writes the whole set. The threshold mask here takes in almost every voxel of the 192 by 256 by 256 volume, which is why its volume is 12581559 mm<sup>3</sup> and its sphericity 0.800, close to the 0.806 of a perfect cube. Sphericity can never exceed 1: a mask that touches the edge of the volume is closed off before its surface is measured. In the full set, entropy is the Shannon entropy of a 64-bin intensity histogram, in bits, and maximum diameter is the largest distance between two points on the mask's surface.

From there, **Save as biomarker** attaches one of the workbench's named, versioned biomarker definitions to the feature set, turning a one-off measurement into a tracked quantity with its own identity.



**Figure 33.** A biomarker is a named, versioned measurement, reviewable and exportable.

Like a segmentation, a biomarker carries its own review status.



**Figure 34.** An accepted biomarker, with the decision recorded beside its value.

### A BIOMARKER CAN ALSO BE A RATIO

Some biomarker definitions, such as a structure's volume normalised by a second structure's, need a second, denominator feature set. The compute control offers a denominator selector only when the chosen definition needs one; a direct measurement (volume, mean intensity, sphericity) does not ask for one at all.

## History

Every action taken on a series is recorded in order, with the parameters it ran with: importing it, each preprocessing step, registration, a manual or AI segmentation, its review, radiomics extraction, a biomarker computed or reviewed. The **History** tab sits next to **Details** in the right-hand panel and reads directly from that record; nothing here is inferred after the fact from whatever a series' current state happens to be. It updates the moment an action completes.

| Details                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | History |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|
| <b>1. Imported</b> 9/22/2026, 11:05:00 AM<br><pre>{   "series_count": 1,   "deidentify_requested": true }</pre>                                                                                                                                                                                                                                                                                                                                                                                        |         |
| <b>2. Preprocessed</b> 9/22/2026, 11:05:25 AM<br><pre>{   "source_series_id": "4234f2b394",   "output_series_id": "77daa9c7ee",   "pipeline_version": "8800b3f8b666",   "steps": [     {       "name": "reorient_canonical",       "params": {         "applied": true       }     },     {       "name": "resample_isotropic",       "params": {         "target_spacing_mm": 1       }     },     {       "name": "gaussian_denoise",       "params": {         "sigma": 1       }     }   ] }</pre> |         |
| <b>3. Segmented (manual ROI)</b> 9/22/2026, 11:05:55 AM<br><pre>{   "method": "threshold",   "source_series_id": "4234f2b394",   "output_series_id": "e0ed1260f9",   "low": 0,   "high": 1000 }</pre>                                                                                                                                                                                                                                                                                                  |         |
| <b>4. Segmented (manual ROI)</b> 9/22/2026, 11:05:58 AM<br><pre>{   "method": "region_growing",   "source_series_id": "4234f2b394",   "output_series_id": "a92bb4621b",   "plane": "axial",   "index": 95,   "seed_voxel": [95, 141, 128],   "tolerance": 50 }</pre>                                                                                                                                                                                                                                   |         |
| <b>5. Segmented (manual ROI)</b> 9/22/2026, 11:06:00 AM<br><pre>{   "method": "manual_rectangle",   "source_series_id": "4234f2b394",   "output_series_id": "7fd7a6154",   "plane": "axial",   "index": 95 }</pre>                                                                                                                                                                                                                                                                                     |         |
| <b>6. Segmented (manual ROI)</b> 9/22/2026, 11:06:04 AM<br><pre>{   "method": "manual_polygon",   "source_series_id": "4234f2b394",   "output_series_id": "e4e816aa27",   "plane": "axial",   "index": 95,   "points": [     {       "x": 107.97948717948718,       "y": 89.533333     }   ] }</pre>                                                                                                                                                                                                   |         |


**Figure 35.** Every action on this series, in order, with its own parameters: import, preprocessing, review and biomarkers alike.

This is the record the product's claim of full provenance rests on. A series' **Description** in the Details tab names what produced it in short form; History is where the complete, ordered account lives.


# AI models

## The model catalogue

Every admitted model appears in one list, disabled for **Run** until a series of a modality it supports is loaded. The workbench never lets you attempt a mismatch, and the note at the top of the dialog names the loaded series' modality so it is clear why some models are greyed out.



**Figure 36.** Every admitted model: task, modality and an estimated CPU runtime.




**Figure 37.** A runtime other than plain TorchScript is called out explicitly, not hidden.

Two things are stated for every model. Its **task**, which decides where its result lands: a segmentation becomes a mask in the ROI list, a detection becomes boxes and a count, a classification becomes a ranked list. And its **runtime**. All eleven run as TorchScript; six of them (HD-BET, prostate, whole-body CT, lung nodule, vertebrae and hippocampus) use `torchscript_sliding_window`, where the network itself is TorchScript but a real orchestration library, nnU-Net's or MONAI's own sliding-window inferer, drives the tiling. The panel says so plainly rather than pretending every model follows one code path. Appendix C is the full catalogue, with what each model returns and its measured accuracy.

A model runs as a background job, listed under the catalogue with its status. While it is queued the row says how many jobs are ahead of it, and if the workbench ever loses track of a job's status it says so on the row instead of spinning indefinitely. A job that was still running when the workbench was restarted is reported as failed, with a message saying the restart interrupted it and that it should be run again.


## Running a segmentation model

HD-BET is brain extraction: an nnU-Net network converted to TorchScript and verified to match its original output exactly. It returns one label, brain, over the whole volume. Running it on the NFBS scan produces a new segmentation series and, because one now exists, switches on the AI overlay.




**Figure 38.** A new segmentation appears as an overlay automatically, at an adjustable opacity.

A legend line under the overlay selector names the mask's labels, so an overlay is never an unexplained colour.




**Figure 39.** HD-BET brain extraction, overlaid on the source volume.



**Figure 40.** A model result reviews exactly like a manual one.

The model's own result sits in the same ROI/Segmentation list as the five manual tools, reviewed the same way, exported the same way.




**Figure 41.** Measure: the voxel count, the volume in cubic millimetres from the voxel spacing, and the mean intensity under the mask.


**Measure** on the row gives the three numbers a mask is most often wanted for, without going through radiomics: how many voxels it holds, the volume they make in cubic millimetres using the series' own spacing, and the mean intensity of the source image under the mask. For a brain mask that volume is the intracranial volume that a hundred neuroimaging papers normalise by.

## Benchmarking a model

Two of the eleven models ship with a real benchmark dataset the workbench can run against on demand: not a single test image, but the model's full reference set, scored the way its own admission was validated.



**Figure 42.** A benchmark runs the model against its own reference dataset, not a single image.




**Figure 43.** Every benchmark run, with the dataset and how it was validated.

HD-BET's most recent run: **Dice 0.9835** (min 0.9824, max 0.9864) against five NFBS subjects' manually verified ground truth. The blood-cell detector: **precision 0.768, recall 0.684** against 87 held-out BCCD images, IoU of 0.5 or better required for a match. Both numbers are exactly what the panel reports. Nothing is rounded up, and nothing is cherry-picked from a better run.

# Detections and classification

## Detection counts and the class filter

A detection model returns boxes, one per finding, each with a class and a confidence. Running the blood-cell detector on a smear draws a box per cell and, critically, counts them.



**Figure 44.** Every red cell found and boxed.

A real blood smear can carry dozens or hundreds of individual cells across several classes: red cells, white cells, platelets. Counting boxes by eye was never a realistic way to read that result, so the sidebar totals them automatically, by class. There is one row here, since this particular field only returned red cells above the model's own confidence threshold.


| CLASS | COUNT |
|-------|-------|
| RBC   | 12    |
| Total | 12    |

**Figure 45.** Every detection counted automatically, by class.

A checkbox per class filters which classes draw at all, on both the image and the list beneath it: one checkbox on this field, more on a smear with a mixed population.



**Figure 46.** A checkbox per class, to isolate one on the image itself.



**Figure 47.** Unchecked: the class is hidden, from both the list and the image.

### THIS IS EXACTLY WHAT THE BENCHMARK NUMBER SAYS TO EXPECT

The plainly visible white cell near the centre of this field is not boxed at all. That is a real miss, not a display artefact. Appendix C's benchmark figure for this model is precision 0.768, recall 0.684 against BCCD's own held-out set: roughly one real cell in three goes undetected. A count panel makes that kind of gap visible at a glance rather than requiring you to eyeball a busy field for what is missing.

## Reviewing detections

Any individual detection can be rejected. It is struck through in the list and excluded from the count, without deleting the underlying record.

|           |         |
|-----------|---------|
| RBC (99%) | Restore |
| RBC (98%) | Reject  |
| RBC (97%) | Reject  |
| RBC (97%) | Reject  |
| RBC (97%) | Reject  |
| RBC (96%) | Reject  |
| RBC (95%) | Reject  |
| RBC (94%) | Reject  |
| RBC (90%) | Reject  |
| RBC (87%) | Reject  |
| RBC (65%) | Reject  |
| RBC (55%) | Reject  |

**Figure 48.** A single detection rejected: struck through, and out of the count.

## Classification and its review

A classification model returns a ranked list of class predictions rather than a box. The histopathology metastasis classifier is shown here, on an H&E-stained lymph node patch, returning tumour against normal with a probability for each.

| Classification             |          |
|----------------------------|----------|
| <b>Metastasis (100.0%)</b> |          |
| Metastasis                 | 100.0%   |
| Other                      | 0.0%     |
| Confirm                    | Override |

**Figure 49.** A classification model returns a ranked list of predictions, not a box.

**Confirm** or **Override** records a human decision against the model's own prediction, the same review discipline as a segmentation or a detection.

**Classification**

**Metastasis (100.0%)**

confirmed

|            |        |
|------------|--------|
| Metastasis | 100.0% |
| Other      | 0.0%   |

Confirm

Override

**Figure 50.** Confirmed: the same review trail as a segmentation or a detection.


### TRUE COLOUR MATTERS FOR THIS PARTICULAR MODEL

Histopathology classification depends on stain colour, which the workbench's shared grayscale preview path would otherwise discard. This model is one of the few that reads the original colour image directly for inference rather than the grayscale-collapsed view everything else shares. Testing confirmed why: none of two metastasis cases was classified correctly through the grayscale path, and all four through the true-colour one.

## Two more modalities

### X-ray classification

Two independent models read chest and wrist X-rays respectively: a Vision Transformer that returns a probability of pneumonia against normal, and SigLIP2, which returns fractured against not fractured.



**Figure 51.** A Vision Transformer classifying a chest X-ray.




**Figure 52.** A second, independent model reading the same modality.

The fracture classifier is the one model in the catalogue whose accuracy on this guide's own independent sample is stated as a limitation: it read two of four wrist images correctly on a dataset it was not trained on. The catalogue entry says so, because a model's result is only as useful as the honesty of the number beside it.

## Parasite detection on a blood smear

A YOLOv8 detector reads a Giemsa-stained smear and returns a box per parasite. It is the fastest model in the catalogue, at a few hundredths of a second per field, and it was spot-checked here against 400 real, balanced images from the standard NIH malaria set rather than trusted on its published figure alone.



**Figure 53.** Parasite detection on a Giemsa-stained blood smear.

Like every detection model, its boxes land in the Detections panel with a count per class, and each can be rejected on review.


## Longitudinal tracking

Repeated scans for one participant, grouped by the reference given at import, can be measured against each other rather than only viewed side by side. Every study imported under the same participant reference is one visit, and the label on each visit card can be edited to name it.

The screenshot shows the 'Longitudinal Imaging' window. At the top, a note states: 'Repeated scans are grouped by the participant reference given at import. Change over time is presented as numbers only — this app never labels a change as clinical improvement or deterioration.' Below this, the 'PARTICIPANT' section shows the ID 'A00061709'. The first visit, 'Baseline', is dated '9/17/2026' and has an 'IMAGE SERIES' of 'sub-A00061709\_ses-NFB3\_T1w\_RAS.nii.gz (source)' and a 'MASK (FOR VOLUME)' of 'threshold'. The second visit, 'Visit 2', is also dated '9/17/2026' and has the same image series and mask. A 'Register to baseline' button is visible next to the second visit's mask. At the bottom, there are buttons for 'Export to Analytics Workbench', 'Export to STUDYPRO', 'Compute Trajectory', and 'Close'.

**Figure 54.** Two visits of one participant, each with an image series and a mask to compare.

Each visit card chooses which image series and which mask represent that visit, independently, so a visit can use its own AI segmentation while another uses a manually drawn one. Every mask of the visit is offered, with its review status in the ROI list to guide the choice. **Register to baseline** on a later visit runs the registration of the previous chapter with the baseline as the fixed series, and adds the registered series to that visit. A mask must match its visit's image in dimensions and voxel spacing; one that does not, such as a mask drawn on a resampled copy, is refused, with the visit named, when the trajectory is computed, because the volume it gave would be ambiguous.



**Figure 55.** A measured volume change: numbers only, never a clinical claim.

The trajectory is exactly that: a measured absolute and percentage change in volume between visits, one point per visit. It is never labelled as improvement or deterioration. That judgement belongs to a clinician reading it in context, not to this software.

### WHAT THE PAIR IN THIS GUIDE IS

The two visits shown are the same NFBS scan imported twice under one participant reference, each with a threshold mask drawn at different bounds. The change the trajectory reports is therefore the difference between two masks, not between two scans. A real pair of visits is chosen, registered and measured in exactly the same way, and would show the same kind of chart with a real interval between the points.

## Exporting a trajectory

Two export routes, both from the same panel, both writing to this machine's own disk.

Exported 2 visit(s), 1 biomarker column(s). [Download CSV](#) - [Download sidecar JSON](#)

**Figure 56.** Exported as a tidy table the Analytics Workbench can read directly.

Exported STUDYPRO biomarker template. [Download .xlsx](#)

Import via Study → Download/Import Template in STUDYPRO — biomarker names must be enabled in that study's settings first.

A00061709/2026-09-17: more than one "Structure volume" value — kept the most recently computed one.

A00061709/2026-09-17: more than one "Structure volume" value — kept the most recently computed one.

A00061709/2026-09-17: more than one "Structure volume" value — kept the most recently computed one.

A00061709: more than one visit was imported at 2026-09-17 02:54; STUDYPRO will see them as the same time point.

A00061709/2026-09-17: more than one "Structure volume" value — kept the most recently computed one.

A00061709/2026-09-17: more than one "Structure volume" value — kept the most recently computed one.

A00061709/2026-09-17: more than one "Structure volume" value — kept the most recently computed one.

**Figure 57.** Or as a biomarker template ready for import into a STUDYPRO study.

The Analytics export is a tidy CSV plus a sidecar manifest, readable by the Analytics Discovery Workbench's own batch import without reshaping. The STUDYPRO export is a biomarker-template workbook, imported through that product's Study, Download and Import Template workflow; biomarker names have to be enabled in that study's settings first. Both of these are one participant at a time.

The two visits here were imported on the same day, in the same minute, and neither card has been named. The Analytics export gives the second visit its own session, `2026-09-17-2`, so the two rows never merge into one, and records a warning saying so in the sidecar. The STUDYPRO template keys a visit by its date and time, so it cannot keep them apart; the warnings under the STUDYPRO export say so, in the line that says STUDYPRO will see them as the same time point. The other lines are the batch of the next chapter at work: each run adds another structure volume to both visits, and the export keeps the most recent and names the rest. Naming each visit on its card avoids the collision in both formats. The next chapter's cohort export writes the same two formats across every participant at once.

## Recipes and batch

Everything so far has been one series at a time: load it, preprocess it, run a model, extract, save a biomarker. A study with fifty participants would mean doing that fifty times. **Recipe & Batch**, the sixth icon on the rail, turns the sequence into a recipe and replays it over any set of studies as one background job.

**Recipe & Batch**

A recipe records what to ask of every study: preprocessing steps, AI models, radiomics and the biomarkers to save. A batch replays it over the studies you tick, one source series at a time, and reports for each what ran, what was skipped by name and why, and what failed. Nothing already in a study is overwritten; every result is a new derivative.

**Recipe**

**PREPROCESSING STEPS (IN THIS ORDER)**

- ☐ Reorient to canonical (RAS+)
- ☐ Resample to isotropic spacing target\_spacing\_mm:
- ☐ N4 bias field correction shrink\_factor:
- ☐ Gaussian denoise sigma:
- ☐ Intensity normalisation mode:
- ☐ Otsu foreground mask

**AI MODELS**

- ☐ Faster R-CNN (MobileNetV3-FPN) — Blood Cell Detection (detection, microscopy, ~0.33s per series)
- ☐ HD-BET Brain Extraction (segmentation, mri, ~133s per series)
- ☐ ONNX Conversion Pipeline Validation (MNIST classifier) (classification, diagnostic\_test, ~0.01s per series)
- ☐ ViT Chest X-ray Pneumonia Classification (classification, xray, ~0.13s per series)
- ☐ SigLIP2 Bone Fracture Classification (classification, xray, ~0.15s per series)
- ☐ ResNet50 Lymph Node Metastasis Classification (PatchCamelyon) (classification, microscopy, ~0.01s per series)
- ☐ MONAI Prostate MRI Zonal Segmentation (segmentation, mri, ~8s per series)
- ☐ MONAI Whole-Body CT Multi-Organ Segmentation (104 structures) (segmentation, ct, ~45s per series)

**Studies** [Select all](#)

- ☐ DEMO-MRI-01 mri · 1 source series · 9/22/2026
- ☐ A00061709 mri · 1 source series · 9/17/2026
- ☐ A00061709 mri · 1 source series · 9/17/2026

[Run recipe on selected studies](#) [Close](#)

**Figure 58.** A recipe on the left, the studies to run it on on the right.

A recipe records **what to ask for**, never what came back: which preprocessing steps and with which parameters, which models, whether to extract radiomics from the segmentations those models produce, and which biomarkers to save from each feature set. The recipe in this chapter is a Gaussian denoise, then HD-BET, then radiomics and the structure volume biomarker, over the guide's own NFBS study and both visits of the longitudinal participant.

**Recipe & Batch**

series)

- ☐ VIT Chest X-ray Pneumonia Classification (classification, xray, ~0.13s per series)
- ☐ SigLIP2 Bone Fracture Classification (classification, xray, ~0.15s per series)
- ☐ ResNet50 Lymph Node Metastasis Classification (PatchCamelyon) (classification, microscopy, ~0.01s per series)
- ☐ MONAI Prostate MRI Zonal Segmentation (segmentation, mri, ~8s per series)
- ☐ MONAI Whole-Body CT Multi-Organ Segmentation (104 structures) (segmentation, ct, ~45s per series)
- ☐ MONAI 3D RetinaNet Lung Nodule Detection (detection, ct, ~90s per series)
- ☐ nnU-Net Vertebrae Segmentation (VerSe) (segmentation, ct, ~60s per series)
- ☐ nnU-Net Hippocampus Segmentation (MSD Task04) (segmentation, mri, ~5s per series)
- ☐ YOLOv8 Malaria Parasite Detection (detection, microscopy, ~0.04s per series)

**RADIOMICS AND BIOMARKERS**

- ☒ Extract radiomics from every segmentation the models produce bins
- ☒ Save Structure volume (mm<sup>3</sup>)
- ☐ Save Mean intensity (intensity units)
- ☐ Save Sphericity (dimensionless)

- Preprocess: Gaussian denoise (sigma = 1.0).
- Run HD-BET Brain Extraction (segmentation); series whose modality is not mri are skipped by name.
- Extract radiomics from every segmentation the models produce, 32 discretisation bins.
- Save Structure volume (v1.0, mm<sup>3</sup>) from each feature set.

**Cohort export**

One table across every participant with biomarkers, one row per visit.

**Figure 59.** Checked: the recipe read back in words, including what will be skipped and why.


**Check recipe** reads the recipe back as sentences before anything runs, and refuses one that cannot work: an unknown step, a step parameter outside the ranges the Preprocess dialog enforces, radiomics with no segmentation model to feed it, or a ratio biomarker, which needs a denominator no batch can choose for you. Ratios stay a per-study action. A recipe loaded from a file is checked the same way.

**Download recipe** writes the recipe as a small JSON file, and **Load recipe** reads one back and checks it, so a group's agreed pipeline is a file rather than a memory. The recipe above, downloaded, is this.

```
{
  "version": "1.0",
  "preprocess": [
    {"name": "gaussian_denoise", "params": {"sigma": 1}}
  ],
  "models": ["hdbet_v1"],
  "radiomics": {"discretization_bins": 32},
  "biomarkers": ["structure_volume"]
}
```

## SKIPPED BY NAME, NOT FAILED

A recipe is applied to every source series in every ticked study, and not every step suits every series. A model that does not accept a series' modality, HD-BET on a CT, say, is skipped for that series, and the report says which step, which series and why. That is the honest answer to "run this over everything": *these steps did not apply*, rather than an error that stops the run.



**Figure 60.** One background job, with progress per study.

The batch is a single job. The row shows its queue position while it waits for a worker and a progress bar with the study it is on while it runs. You can keep working, or close the dialog and come back. This one ran HD-BET three times, about two minutes each.

| Batch 3130cd525407 — started 9/22/2026, 11:08:22 AM |                                              |                                                                                                                      |                 |       |
|-----------------------------------------------------|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------|-----------------|-------|
| PARTICIPANT                                         | SERIES                                       | RAN                                                                                                                  | SKIPPED BY NAME | ERROR |
| DEMO-MRI-01                                         | sub-A00032875_ses-NFB3_T1w.nii.gz<br>mri     | preprocessed (83d5e980267b)<br>hdbet_v1: segmentation<br>radiomics<br>Structure volume = 115523.381 mm <sup>3</sup>  |                 |       |
| A00061709                                           | sub-A00061709_ses-NFB3_T1w_RAS.nii.gz<br>mri | preprocessed (83d5e980267b)<br>hdbet_v1: segmentation<br>radiomics<br>Structure volume = 1318263.239 mm <sup>3</sup> |                 |       |
| A00061709                                           | sub-A00061709_ses-NFB3_T1w_RAS.nii.gz<br>mri | preprocessed (83d5e980267b)<br>hdbet_v1: segmentation<br>radiomics<br>Structure volume = 1318263.239 mm <sup>3</sup> |                 |       |

**Figure 61.** Per series: what ran, what was skipped by name, what failed.

When it finishes, the results table has one row per series: what ran (the new series, the model result, the feature set, each biomarker's value), what was skipped and why, and any error. A failure in one series stops that series only; the batch carries on with the rest. Every derived series, feature set and biomarker it produced sits in its study exactly as if it had been made by hand, reviewable, with its full History.

|                                                          |      |
|----------------------------------------------------------|------|
| 9/22/2026, 11:08:22 AM — 3 studies, 3 series, 0 error(s) | View |
| 9/18/2026, 7:09:35 PM — 3 studies, 3 series, 0 error(s)  | View |

**Figure 62.** Every batch that has run, kept with its study and series counts.

Every run is kept and listed under **Previous batch runs**, with the date, the number of studies and series, and the error count, and **View** brings its results table back.

## Cohort export

Cohort export

One table across every participant with biomarkers, one row per visit.

[Export cohort to Analytics Workbench](#)
[Export cohort to STUDYPRO](#)

Exported 3 visit(s), 1 biomarker column(s). [Download CSV](#) · [Download sidecar JSON](#)


Skipped: DEMO-PROSTATE-01 (None of these studies have any computed biomarkers yet — compute at least one before exporting.); TEST-001 (None of these studies have any computed biomarkers yet — compute at least one before exporting.)

**Figure 63.** Every participant with biomarkers, one table.


The two buttons at the foot of the dialog write the same two formats the Longitudinal panel does, but across **every** participant: one row per visit, one column per biomarker, for the Analytics Workbench as a CSV with its sidecar manifest, or for STUDYPRO as a biomarker-template workbook. A participant with no biomarkers yet is named under the message and in the sidecar's `skipped` list rather than silently left out, so the table's row count is never a surprise. Here those are the prostate study and the DICOM study of the chapters that follow, which carry no biomarkers.

## A second anatomy model

The workbench is not a single-organ tool. Prostate MRI zonal segmentation, a MONAI sliding-window inferer independently verified against real Prostate158 test cases, runs the same way HD-BET does, on a completely different scan and returns two structures rather than one: the central gland and the peripheral zone. Measured on the three Prostate158 cases it ships with, the model scores a mean Dice of 0.85 across both zones, inside the range its authors publish.



**Figure 64.** A second, independent anatomy model: the central gland and the peripheral zone of the prostate.



**Figure 65.** A multi-structure mask names every label it carries under the overlay selector.

The legend under the overlay selector now lists two labels with their codes, which is what lets a two-structure mask be read at all. Radiomics and Measure in the sidebar work on the whole mask, and a biomarker such as a zone volume ratio is where the denominator selector described earlier comes in. A feature set records which label it measured: a single-label mask records its label, and

the union of a two-label mask records none, so the two zones together are never passed off as one structure. The workbench's radiomics request can also be told one label to measure, and refuses a label the mask does not contain.

Appendix C lists the four models this guide does not walk through, all on CT: whole-body multi-organ segmentation, lung nodule detection, vertebrae segmentation and hippocampus segmentation on MRI. Each was admitted and verified the same way as the ones shown here, and each returns its result into the same panels.

## DICOM and de-identification

DICOM is the one format where the workbench does not have to infer orientation from array axis position. The file's own per-slice geometry is read directly, and the orientation it gives is recorded with the series, so *Reorient to canonical* works on it and the models that expect a volume in RAS orientation receive the scan the right way round. A multi-frame file becomes one volume, a MONOCHROME1 image, stored with white as zero, is inverted so that it reads like every other series, and a colour DICOM becomes grayscale.

**Import Imaging Data**

**PARTICIPANT REFERENCE (OPTIONAL)**

TEST-001

Free text only — this app does not connect to STUDYPRO directly.

**MODALITY**

CT

Used when the source format does not state a modality (NIFTI, plain images). DICOM overrides this with its own Modality tag where recognised.

☒ **De-identify DICOM series on import**

Writes a separate de-identified derivative; the original file is never modified. This does not guarantee every identifying element is removed.

**FILES**

Drag and drop files here or click to browse

DICOM (.dcm), NIFTI (.nii, .nii.gz), images (.png/.jpg/.tif), or a .zip of any of these

CT\_small.dcm (38.3 KB)

Import Cancel

**Figure 66.** DICOM carries its own Modality tag, and de-identification writes a separate derivative rather than changing the original file.




**Figure 67.** A CT slice imported from real DICOM, de-identified on the way in.

De-identification, when checked, writes a separate de-identified copy alongside the original rather than modifying what you uploaded. The original stays exactly as imported, in case the de-identified derivative ever needs to be checked against it. Everything the viewer and the tools then do with that series reads from the de-identified copy; the original is kept, not used.


The copy has every patient and institution identifier replaced (name, ID, birth date, address, accession number, physicians, operator, station and department), every date and time blanked, every comment field emptied, study, procedure, admission, order and trial identifiers, the additional patient history and the device serial number blanked, and every private tag removed. The same rules are applied inside nested sequences, where an identifier is removed outright. PatientAge, PatientSex, PatientSize and PatientWeight are kept, because an analysis usually needs them and none of them names anyone.

# Managing studies



**Figure 68.** Every dataset used in this guide, as one catalogue.

Deleting a study removes it and everything derived from it, every preprocessed series, every segmentation, every biomarker, and asks first. If the study you delete is the one on screen, the viewer, the compare pane and every sidebar panel clear with it.



**Figure 69.** Deleting a study is confirmed, and removes its derived data too.

## THERE IS NO UNDO

Confirming a deletion removes the study's files from disk. There is no recycle bin and no soft-delete. Re-importing the source data is the only way back.

## Appendix A: Control reference

| Control                        | Where                                                      | Does                                                                              |
|--------------------------------|------------------------------------------------------------|-----------------------------------------------------------------------------------|
| Import                         | Sidebar                                                    | Bring in DICOM, NIfTI or a plain image                                            |
| Studies                        | Sidebar                                                    | Browse, view and delete studies                                                   |
| Preprocess & Register          | Sidebar                                                    | Run preprocessing steps or register two series                                    |
| AI Models                      | Sidebar                                                    | Run or benchmark an admitted model                                                |
| Longitudinal                   | Sidebar                                                    | Track a participant across repeated visits                                        |
| Recipe & Batch                 | Sidebar                                                    | Build, check, save and load a recipe; run it over many studies; export the cohort |
| About                          | Sidebar                                                    | Version and build                                                                 |
| Axial / Coronal / Sagittal     | Viewer toolbar                                             | Switch the displayed plane                                                        |
| Zoom in / out / reset          | Viewer toolbar                                             | Magnification, not a measurement                                                  |
| Measure                        | Viewer toolbar                                             | Two-click distance, in real millimetres                                           |
| Compare                        | Viewer toolbar                                             | Open a second series alongside the first, with its own plane and slice controls   |
| Window Width / Center          | Sidebar                                                    | Display contrast only, never written to the series                                |
| ROI tool selector              | Sidebar                                                    | Threshold, region growing, rectangle, polygon, brush                              |
| View / Measure / Radiomics     | Per segmentation row                                       | Show the mask; voxel count, volume and mean intensity; the full feature set       |
| Accept / Reject                | Segmentation, detection, biomarker and classification rows | Record a review decision in provenance                                            |
| Save as biomarker              | Radiomics result                                           | Attach a named, versioned biomarker definition                                    |
| Run                            | Per model                                                  | Apply that model to the loaded series                                             |
| Benchmark                      | Per model, where available                                 | Score the model against its own reference dataset                                 |
| Check / Download / Load recipe | Recipe & Batch                                             | Read a recipe back in words; save or restore it as JSON                           |
| Run recipe on selected studies | Recipe & Batch                                             | Replay the recipe over the ticked studies as one job                              |

| Control             | Where          | Does                                                                                 |
|---------------------|----------------|--------------------------------------------------------------------------------------|
| Previous batch runs | Recipe & Batch | Every earlier run, with View to bring its table back                                 |
| Export cohort       | Recipe & Batch | Every participant's biomarkers in one table, for the Analytics Workbench or STUDYPRO |

## Appendix B: Supported formats and modalities

| Format                                             | Notes                                                                                                                                                                                             |
|----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DICOM ( <code>.dcm</code> and extension-less)      | Own orientation and Modality tag read directly and the orientation recorded; multi-frame files load as one volume, MONOCHROME1 is inverted, colour becomes greyscale; de-identification available |
| NIfTI ( <code>.nii</code> , <code>.nii.gz</code> ) | Modality must be given at import; axis order assumed canonical, and checked at import; 3-D and 4-D only, a 2-D file is refused                                                                    |
| Plain image (PNG, JPEG, TIFF)                      | Read as microscopy or X-ray per the modality given at import; 16-bit images keep their full range                                                                                                 |

Modalities the workbench recognises: **mri**, **ct**, **xray**, **microscopy**.

## Appendix C: The model catalogue

| Model                                                         | Task           | Modality   | Returns                                                                       | CPU time |
|---------------------------------------------------------------|----------------|------------|-------------------------------------------------------------------------------|----------|
| HD-BET Brain Extraction                                       | segmentation   | mri        | One label, the brain, over the whole volume                                   | ~133 s   |
| nnU-Net Hippocampus Segmentation (MSD Task04)                 | segmentation   | mri        | Two labels, anterior and posterior hippocampus                                | ~5 s     |
| MONAI Prostate MRI Zonal Segmentation                         | segmentation   | mri        | Two labels: central gland and peripheral zone                                 | ~8 s     |
| MONAI Whole-Body CT Multi-Organ Segmentation                  | segmentation   | ct         | 104 labels, from spleen to urinary bladder, every vertebra and rib among them | ~45 s    |
| nnU-Net Vertebrae Segmentation (VerSe)                        | segmentation   | ct         | One label per vertebra, C1 to L6                                              | ~60 s    |
| MONAI 3D RetinaNet Lung Nodule Detection                      | detection      | ct         | A box per nodule, with confidence                                             | ~90 s    |
| Faster R-CNN Blood Cell Detection                             | detection      | microscopy | A box per cell, classed red cell, white cell or platelet, with counts         | ~0.3 s   |
| YOLOv8 Malaria Parasite Detection                             | detection      | microscopy | A box per parasite                                                            | ~0.04 s  |
| ResNet50 Lymph Node Metastasis Classification (PatchCamelyon) | classification | microscopy | Tumour against normal, with a probability each                                | ~0.01 s  |
| ViT Chest X-ray Pneumonia Classification                      | classification | xray       | Pneumonia against normal, with a probability each                             | ~0.1 s   |
| SigLIP2 Bone Fracture Classification                          | classification | xray       | Fractured against not fractured, with a probability each                      | ~0.2 s   |

Every model was converted to TorchScript and verified against its original framework output before admission. The six sliding-window models rely on nnU-Net's or MONAI's own tiling rather than a hand-rolled equivalent, and the catalogue says so on each. The measured accuracies the guide quotes are HD-BET's Dice of 0.9835 on five NFBS subjects, the blood-cell detector's precision 0.768 and recall 0.684 on 87 held-out BCCD images, the prostate model's mean Dice of 0.85 on three Prostate158 cases, and the fracture classifier's two of four on the guide's own wrist images. The three CT models and the hippocampus model carry their authors' published figures, which the catalogue entry names, and were verified here to reproduce their original framework's output exactly rather than re-scored on new data.

The catalogue also lists an MNIST digit classifier converted through ONNX. It exercises the conversion path and is matched to no imaging modality, so it is never offered for a loaded series.

## Appendix D: Known limitations

---

**NIfTI axis order is assumed, not corrected.** A file whose axis order is permuted rather than merely flipped will show the right data under the wrong plane label until *Reorient to canonical* is applied. The workbench warns about this at import, in *Automatic quality control*, but does not change the file, consistent with never silently modifying an import.

**Registration reports a similarity metric, not a pass or fail score.** The number returned is meaningful for comparing attempts on the same pair of series; it is not a universal threshold for good enough.

**A recipe cannot compute a ratio biomarker, and does not choose series for you.** A ratio needs a denominator feature set that only a person can pick, so it stays a per-study action, and a batch runs on every *source* series of each ticked study, skipping by name whatever a step cannot take, rather than letting you hand-pick one series per study.

**A trajectory or a biomarker is only as trustworthy as the segmentation it was built from.** The workbench records a review decision on every segmentation and biomarker, but it does not filter by it: a rejected mask can still be chosen for a trajectory and still appears in an export, with its status recorded. Choosing what to leave out remains a human decision at every step.

**The three CT models were verified against their original frameworks, not re-scored here.** Whole-body, vertebrae and lung nodule carry their authors' published accuracy; this guide had no CT volume to run them on.

## Appendix E: Data used in this guide

| Data                  | Provenance                                                                                                                                                                        |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NFBS T1               | One raw T1 volume from the Neurofeedback Skull-stripped repository, <code>DemoData/NFBS_Dataset/</code> , whose native axis order is what the quality-control chapter warns about |
| A00061709, two visits | Two canonically oriented T1 volumes of one NFBS participant, present in the study cache as the longitudinal participant                                                           |
| Prostate158           | One T2 volume from the Prostate158 sample, <code>DemoData/Prostate158_Sample/</code>                                                                                              |
| BCCD                  | One blood smear field from the BCCD dataset, <code>DemoData/BCCD_Dataset/</code>                                                                                                  |
| PatchCamelyon         | One metastasis patch, <code>DemoData/PatchCamelyon_Sample/</code>                                                                                                                 |
| Chest X-ray           | One pneumonia image from a Kermany-derived sample, <code>DemoData/ChestXray_Sample/</code>                                                                                        |
| FracAtlas             | One fractured wrist image, <code>DemoData/FracAtlas_Sample/</code>                                                                                                                |
| Malaria               | One parasitised smear from the NIH malaria set, <code>DemoData/Malaria_Sample/</code>                                                                                             |
| DICOM                 | <code>CT_small.dcm</code> , the pydicom project's own test file, so the DICOM chapter needs nothing checked in                                                                    |

None of these is patient data.

# Index

---

## A

---

accepted, 4, 27, 28  
achieved spacing, 19  
AI overlay, 10, 15, 32, 33, 50  
Analytics Discovery Workbench, 3, 5, 44, 45, 49, 56  
axis order, 4, 14, 16, 20, 21, 57, 60, 61

## B

---

BCCD, 6, 35, 36, 37, 41, 42, 58, 61  
benchmark, 35, 37, 55  
bias field, 3, 18  
biomarker, 3, 4, 6, 10, 24, 27, 28, 29, 30, 45, 46, 47, 48, 49, 50, 54, 55, 56, 60  
    ratio, 29, 47, 50, 60  
brush, 3, 25, 27, 55

## C

---

canonical, 14, 16, 18, 19, 20, 21, 52, 57, 60  
classification, 3, 4, 6, 10, 32, 36, 38, 39, 40, 41, 55, 58, 59  
cohort export, 3, 4, 45, 48, 56  
compare pane, 3, 5, 9, 21, 54  
confidence threshold, 36, 58  
Confirm, 11, 38  
coronal, 3, 5, 13, 14, 17, 55

## D

---

de-identification, 6, 11, 52, 53, 57  
deleting a study, 54  
denoise, 18, 19, 46  
detections, 3, 4, 5, 6, 10, 32, 36, 37, 38, 39, 41, 42, 51, 55, 58  
Dice, 35, 50, 58  
DICOM, 3, 4, 5, 6, 11, 17, 21, 49, 52, 53, 55, 57, 61

## E

---

entropy, 28

## F

---

FracAtlas, 6, 40, 41, 58, 61

## H

---

HD-BET, 4, 6, 32, 33, 35, 46, 48, 50, 58  
hippocampus, 32, 51, 58

History, 3, 4, 20, 27, 29, 30, 48, 53, 55, 61

## I

---

Import dialog, 11  
import warnings, 11  
interrupted job, 32  
isotropic, 18, 19

## L

---

legend, 33, 50  
longitudinal, 3, 6, 11, 12, 21, 22, 43, 44, 45, 46, 49, 55, 61

## M

---

malaria, 6, 41, 42, 58, 61  
maximum diameter, 28  
modality, 3, 6, 11, 31, 41, 48, 52, 57, 58, 59  
model catalogue, 3, 4, 5, 31, 32, 41, 54, 55, 58, 59  
MONAI, 32, 50, 58  
multi-frame, 52, 57

## N

---

NFBS, 6, 14, 19, 20, 24, 32, 35, 44, 46, 58, 61  
NIfTI, 3, 4, 5, 11, 14, 21, 55, 57, 60  
nnU-Net, 32, 58

## O

---

Otsu, 18

## P

---

participant reference, 3, 4, 6, 8, 11, 12, 22, 43, 44, 45, 46, 48, 49, 55, 56, 61  
PatchCamelyon, 6, 38, 39, 58, 61  
pipeline version, 20  
plain image, 3, 4, 5, 11, 12, 14, 21, 55, 57  
pneumonia, 6, 40, 58, 61  
polygon, 3, 25, 55  
preprocessing, 3, 4, 6, 12, 14, 15, 18, 21, 29, 30, 46, 55  
previous batch runs, 48, 56  
prostate, 6, 32, 49, 50, 58, 61

## Q

---

quality control, 3, 4, 6, 10, 14, 16, 17, 20, 60

## R

---

radiomics, [3](#), [6](#), [24](#), [26](#), [27](#), [29](#), [35](#), [46](#), [47](#), [50](#), [55](#)  
recipe, [3](#), [4](#), [8](#), [46](#), [47](#), [48](#), [55](#), [56](#), [60](#)  
    batch, [3](#), [4](#), [6](#), [8](#), [32](#), [45](#), [46](#), [47](#), [48](#), [55](#), [56](#), [60](#)  
    download, [45](#), [47](#), [55](#)  
rectangle, [3](#), [24](#), [55](#)  
region growing, [3](#), [24](#), [55](#)  
registration, [3](#), [6](#), [8](#), [12](#), [16](#), [18](#), [21](#), [22](#), [29](#), [43](#), [44](#), [55](#),  
    [60](#)  
rejected, [27](#), [37](#), [38](#), [42](#), [60](#)  
review status, [3](#), [4](#), [27](#), [28](#), [43](#), [55](#), [60](#)  
ruler, [3](#), [15](#), [16](#)

## S

---

sagittal, [3](#), [4](#), [13](#), [14](#), [17](#), [21](#), [55](#)  
same-day visits, [45](#)  
segmentation, [3](#), [4](#), [10](#), [12](#), [18](#), [19](#), [24](#), [25](#), [26](#), [27](#), [28](#),  
    [29](#), [32](#), [33](#), [34](#), [35](#), [38](#), [39](#), [43](#), [44](#), [46](#), [47](#), [50](#), [51](#),  
    [54](#), [55](#), [58](#), [60](#)  
similarity metric, [22](#), [23](#), [60](#)  
skipped by name, [11](#), [47](#), [48](#), [49](#)

slice slider, [3](#), [13](#), [14](#), [15](#), [21](#), [24](#)  
sliding window, [32](#)  
sphericity, [28](#), [29](#)  
Studies dialog, [12](#)  
STUDYPRO, [3](#), [5](#), [11](#), [45](#), [49](#), [56](#)

## T

---

TorchScript, [32](#), [58](#)  
trajectory, [3](#), [27](#), [43](#), [44](#), [60](#)  
true colour, [6](#), [39](#)

## V

---

vertebrae, [32](#), [51](#), [58](#), [60](#)  
voxel count, [35](#), [55](#)

## W

---

whole-body CT, [32](#), [51](#), [58](#)  
window and level, [3](#), [10](#), [55](#)

## Z

---

zoom, [3](#), [15](#), [55](#)