Medical Image Analysis & Radiomics Workshop

From DICOM to Predictive Models

Inês Carvalho - ines.c@edu.ulisboa.pt

The Power of Medical Imaging in Oncology

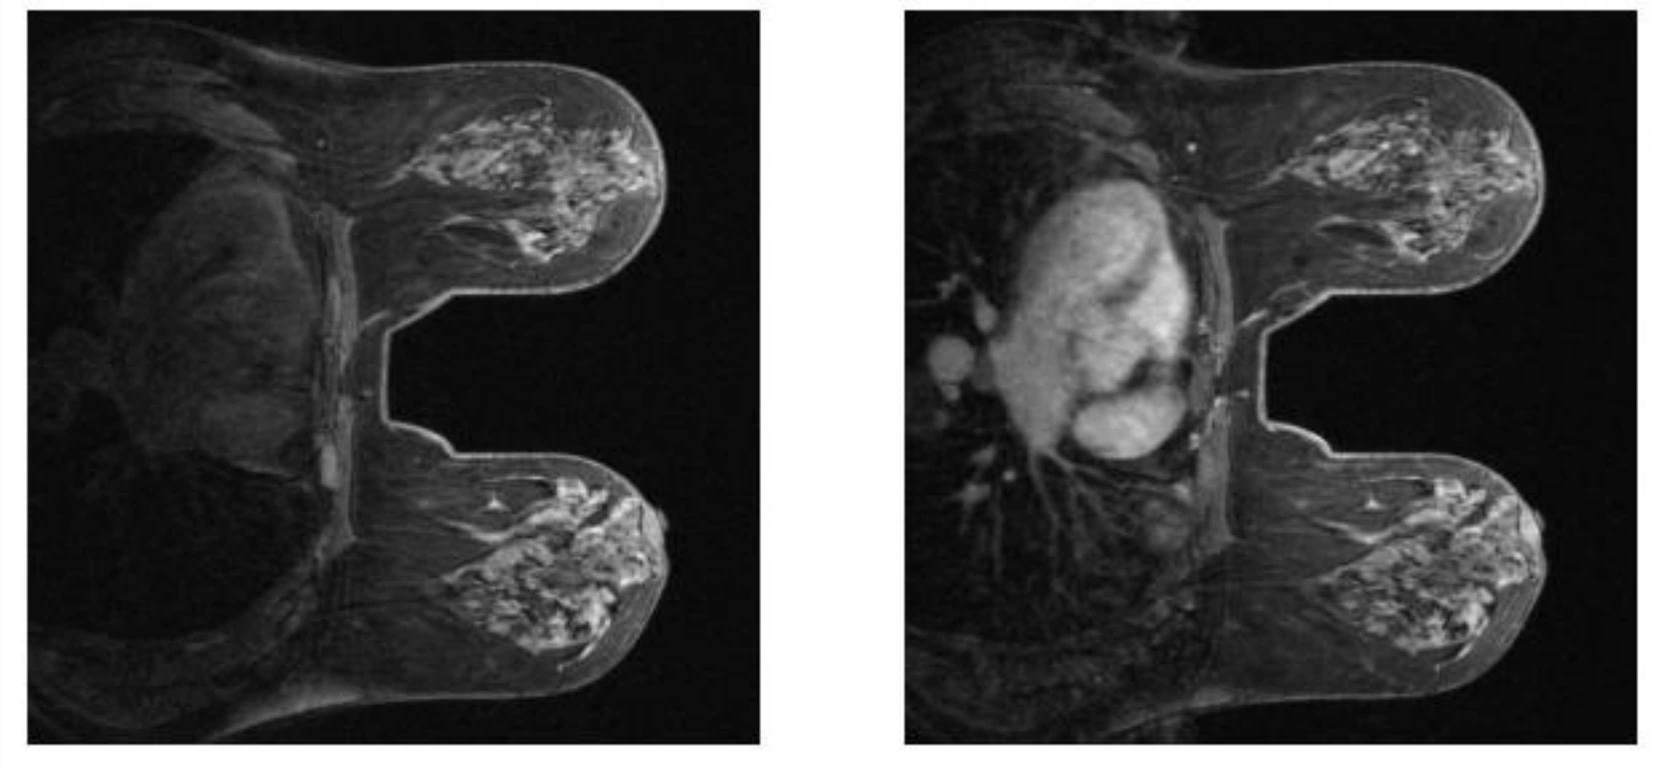

The Power of Medical Imaging in Oncology

Beyond Visual Inspection

AI extracts quantitative patterns invisible to the naked eye — texture, shape, and intensity distributions encode tumour biology.

Mirroring Clinical Perception

Multimodal AI mirrors what clinicians already perceive: combining imaging, pathology, and genomics for holistic diagnosis.

Non-Invasive Biomarkers

Images provide non-invasive, repeatable windows into tumour phenotype, enabling longitudinal monitoring without biopsy.

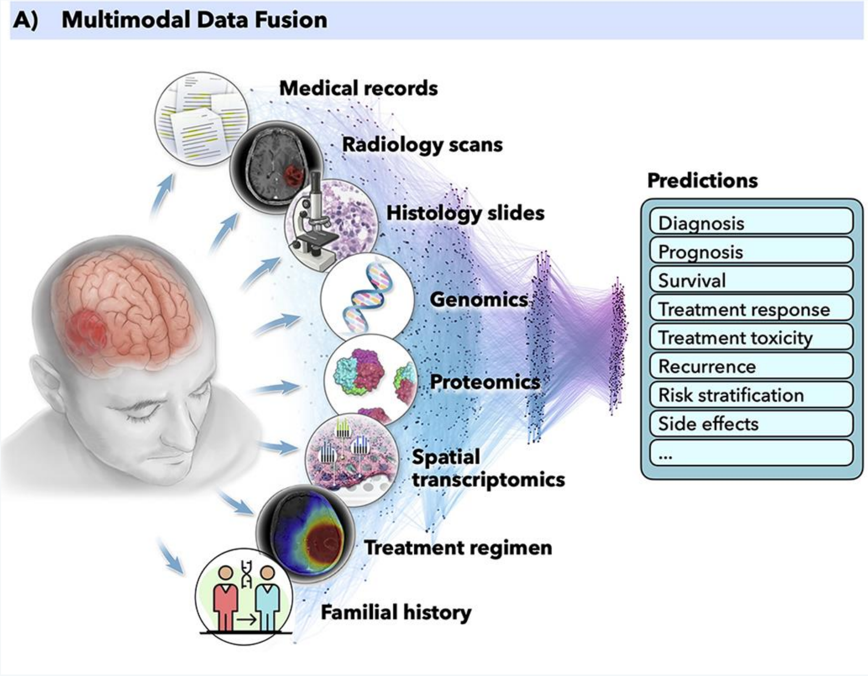

Acosta et al., Nature Medicine 2022 — Multimodal biomedical AI

Imaging Modalities: Diagnostic & Prognostic Tools

Clinical imaging modalities used in oncology — each captures a different data type

X-ray	**Ultrasound**	**CT**	**MRI**	**PET**
Radiography	*US / Echography*	*Computed Tomography*	*Magnetic Resonance*	*Positron Emission Tom.*
2D image	**2D video (temporal)**	**3D volume**	**3D + 4D (DCE / DWI)**	**3D functional**
Bone & lung screening — fast, low-cost first-line exam	Real-time soft tissue imaging — no radiation, bedside use	High-resolution anatomy — tumour staging & lung imaging	Superior soft tissue contrast — brain, breast, oncology	Metabolic activity — detect tumours & metastases

From 2D snapshots to 4D dynamic volumes — all generate quantitative data AI can learn from.

Histopathology Modalities: Tissue-Level Imaging

Histology & pathology modalities used in oncology — from stained slides to spatial-omics

H&E	**IHC**	**mIF**	**Spatial-Omics**	**3D Histology**
Brightfield WSI	*Multiplex IF*	*Multiplex Immunofluorescence*	*Spatial Transcriptomics*	*Confocal / Light-sheet*
2D gigapixel WSI	**2D protein stain**	**2D multi-channel**	**2D + gene expression**	**3D tissue volume**
Gold-standard tissue morphology — diagnosis & grading	Single-marker expression — receptor & subtype profiling	Co-localised markers — tumour microenvironment mapping	Tissue + transcriptome — molecular spatial context	Volumetric architecture — vasculature & 3D tumour analysis

From single-stain slides to spatial-omics volumes — pathology is becoming a high-dimensional, AI-ready modality.

ISTARS Group — Imaging Related Projects

Two clinical imaging projects spanning 2D cephalometric X-ray and 3D brain MRI

Project 1 — Cephalometric X-ray

Skeletal Class III malocclusion subtyping

CLINICAL QUESTION

Can we identify distinct skeletal subtypes of Class III malocclusion from lateral cephalometric X-rays?

APPROACH

- 2D anatomical landmarks on lateral cephalograms
- Generalized Procrustes Analysis (GPA)
- Unsupervised clustering (K-means + UMAP)

Project 2 — Brain MRI

Brain metastasis — surgical decision support

CLINICAL QUESTIONS

Is surgery indicated? What is the expected prognosis after surgery?

APPROACH

- 288 patients — MRI + biopsy + clinical follow-up
- Multi-modal 3D MRI (T1, T1c, T2, FLAIR)
- Radiomics + ML for prognostic modelling

Project 1 — Cephalometric X-ray Analysis

Identifying skeletal subtypes of Class III malocclusion from 2D landmark coordinates

Motivation / Problem Overview

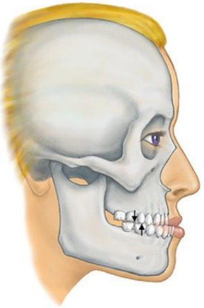

Skeletal Class III
Facial Profile

Cephalometric Measurements

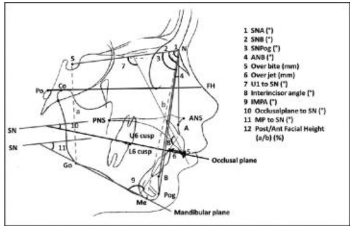

1 SNA (°)
2 SNB (°)
3 SNPog (°)
4 ANB (°)
5 Over bite (mm)
6 Over jet (mm)
7 U1 to SN (°)
8 Interincisor angle (°)
9 IMPA (°)
10 Occlusalplane to SN (°)
11 MP to SN (°)
12 Post/Ant Facial Height (a/b) (%)

2D Landmark Coordinates Data

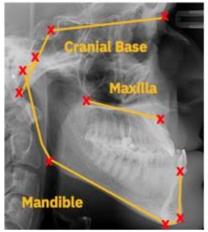

Presented Solution

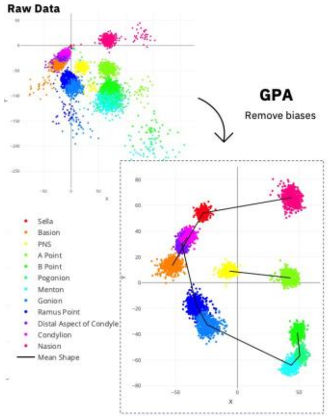

Generalized Procrustes Analysis

Clustering

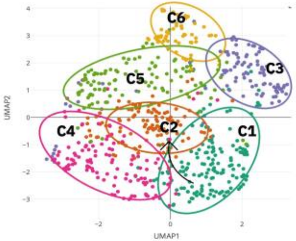

Unsupervised Clustering

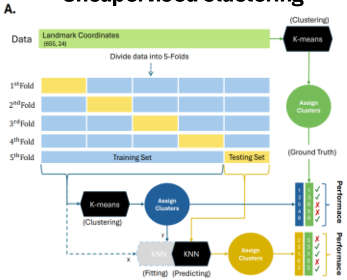

ISTAR Group — Imaging Related Projects

Two clinical imaging projects spanning 2D cephalometric X-ray and 3D brain MRI

Project 1 — Cephalometric X-ray

Skeletal Class III malocclusion subtyping

CLINICAL QUESTION

Can we identify distinct skeletal subtypes of Class III malocclusion from lateral cephalometric X-rays?

APPROACH

- 2D anatomical landmarks on lateral cephalograms
- Generalized Procrustes Analysis (GPA)
- Unsupervised clustering (K-means + UMAP)

Project 2 — Brain MRI

Brain metastasis — surgical decision support

CLINICAL QUESTIONS

Is surgery indicated? What is the expected prognosis after surgery?

APPROACH

- 288 patients — MRI + biopsy + clinical follow-up
- Multi-modal 3D MRI (T1, T1c, T2, FLAIR)
- Radiomics + ML for prognostic modelling

Both projects illustrate how clinical imaging — across data types — can drive quantitative, AI-ready analysis.

Project 2 — Brain MRI · Part 1: Problem & Dataset

Brain metastasis cohort — MRI + biopsy + clinical follow-up to support surgical decision-making

Motivation / Problem Overview

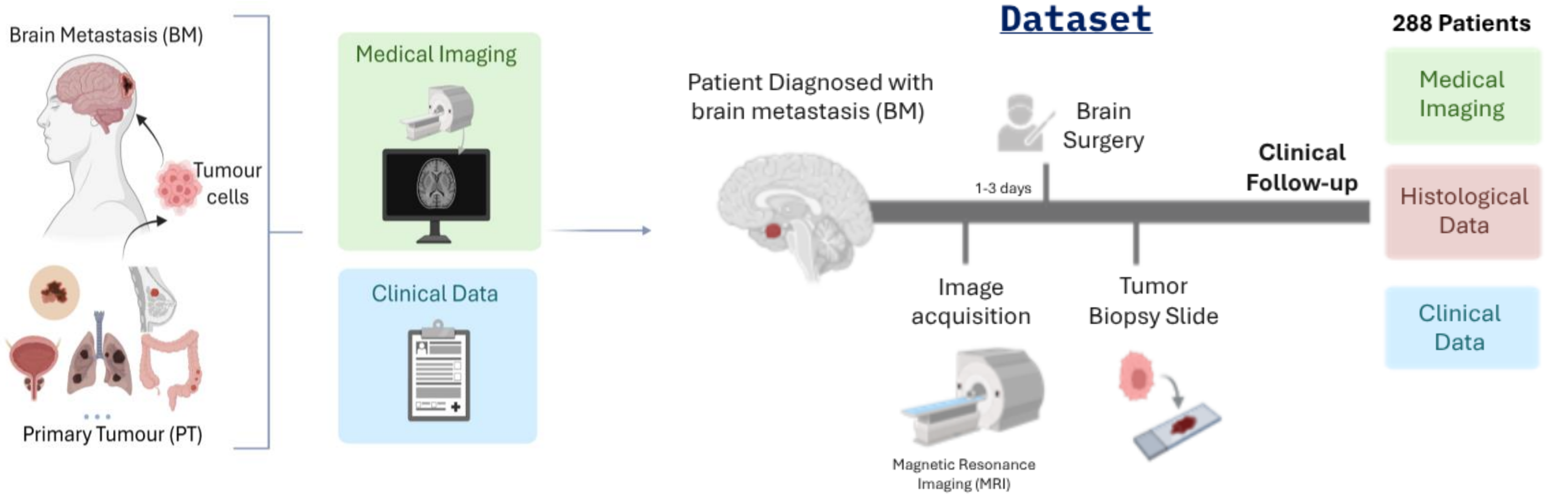

288 patients · multi-modal MRI · histology · longitudinal clinical follow-up

Project 2 — Brain MRI · Part 2: Proposed Solution

Feature-extraction + machine learning pipeline for surgical decision support and prognosis prediction

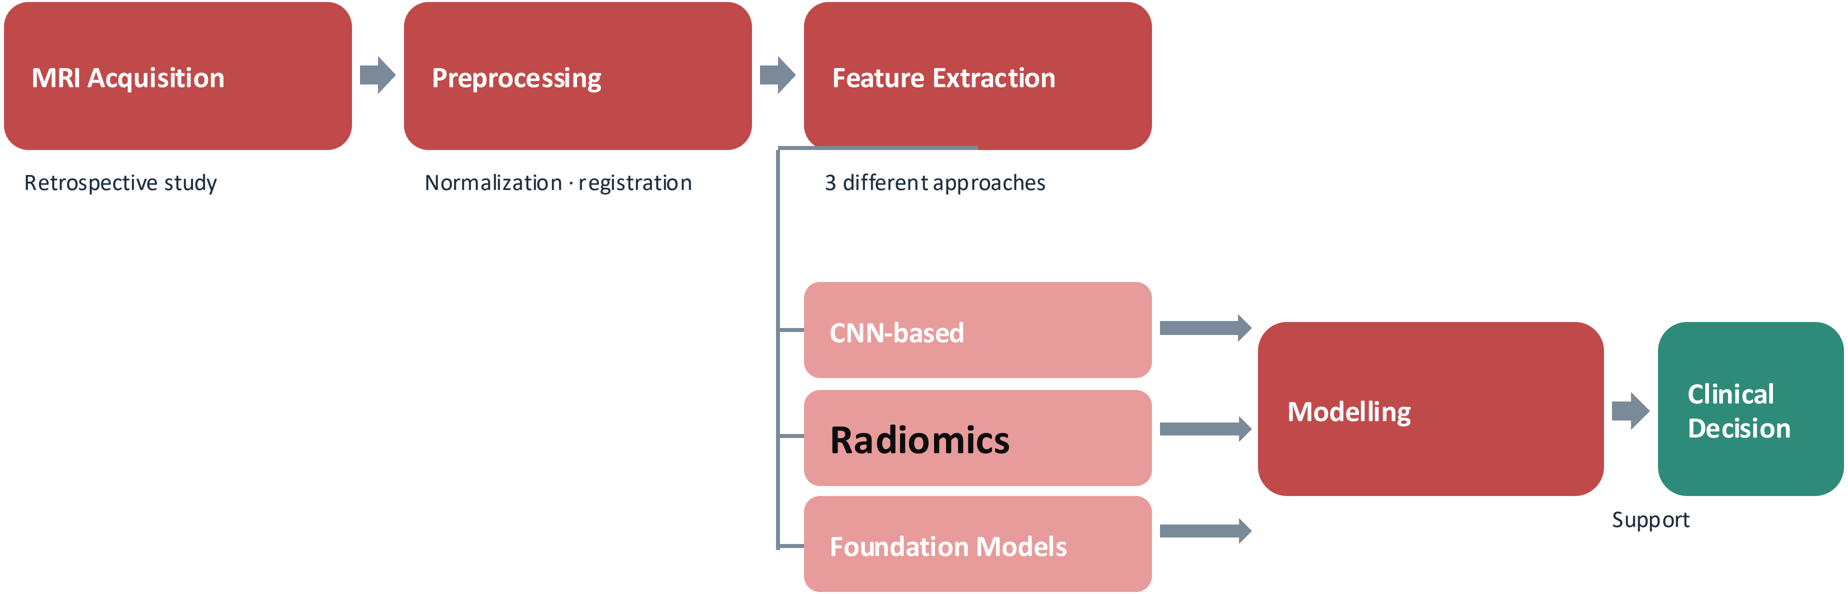

MRI Acquisition

Retrospective study

Preprocessing

Normalization · registration

Feature Extraction

3 different approaches

CNN-based

Radiomics

Foundation Models

Modelling

Clinical Decision

Support

End-to-end pipeline — from raw MRI to support clinical decision

Workshop Overview — Full Day Schedule

S01

Medical Image Formats

DICOM & NIfTI

Hands-on code

Beginner level

S02

MRI Preprocessing

Bias, Denoise, Resample

Hands-on code

Beginner level

S03

Radiomics Features

Shape, Texture, DCE-MRI

Hands-on code

Beginner level

S04

Predictive Modelling

ML on TCGA-BRCA data

Hands-on code

Beginner level

Workshop — Breast MRI (DCE-MRI)

Dynamic contrast-enhanced MRI for breast cancer characterization — TCGA-BRCA cohort

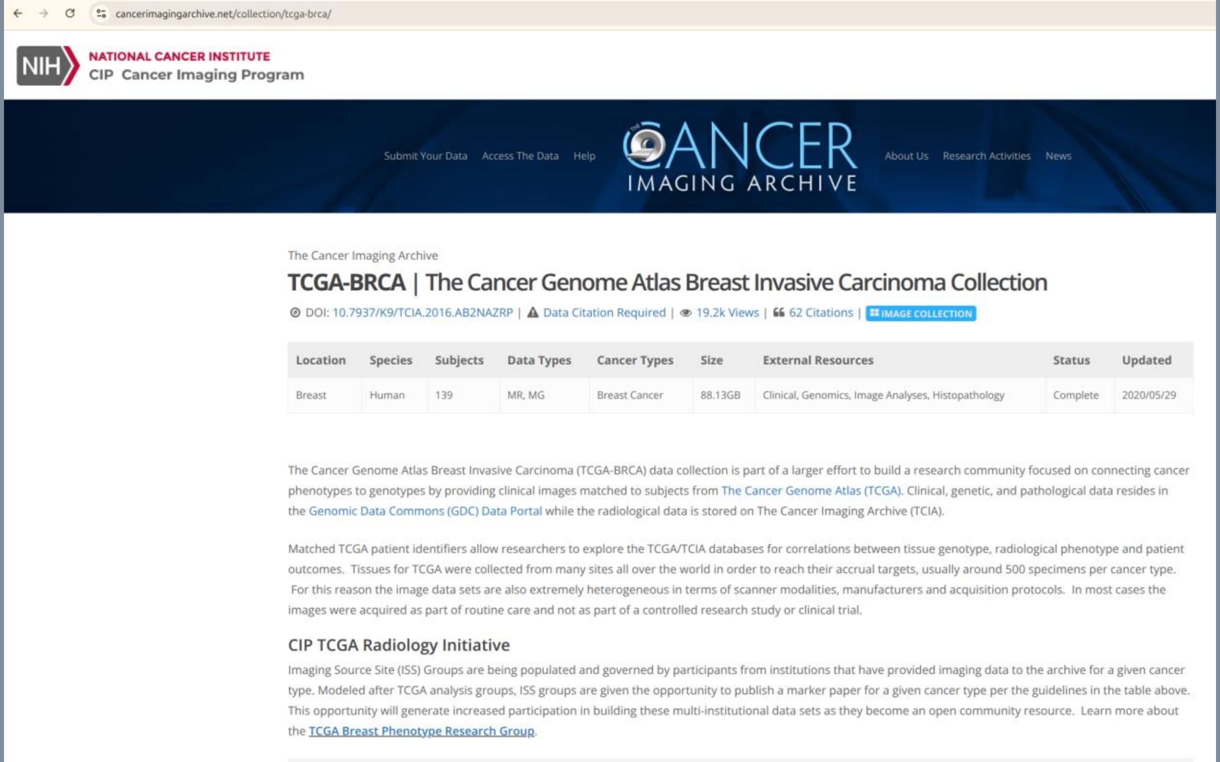

The Cancer Imaging Archive

TCGA-BRCA | The Cancer Genome Atlas Breast Invasive Carcinoma Collection

DOI: 10.7937/K9/TCIA.2016.AB2NAZRP | Data Citation Required | 19.2k Views | 62 Citations | IMAGE COLLECTION

Location	Species	Subjects	Data Types	Cancer Types	Size	External Resources	Status	Updated
Breast	Human	139	MR, MG	Breast Cancer	88.13GB	Clinical, Genomics, Image Analyses, Histopathology	Complete	2020/05/29

The Cancer Genome Atlas Breast Invasive Carcinoma (TCGA-BRCA) data collection is part of a larger effort to build a research community focused on connecting cancer phenotypes to genotypes by providing clinical images matched to subjects from The Cancer Genome Atlas (TCGA). Clinical, genetic, and pathological data resides in the Genomic Data Commons (GDC) Data Portal while the radiological data is stored on The Cancer Imaging Archive (TCIA).

Matched TCGA patient identifiers allow researchers to explore the TCGA/TCIA databases for correlations between tissue genotype, radiological phenotype and patient outcomes. Tissues for TCGA were collected from many sites all over the world in order to reach their accrual targets, usually around 500 specimens per cancer type. For this reason the image data sets are also extremely heterogeneous in terms of scanner modalities, manufacturers and acquisition protocols. In most cases the images were acquired as part of routine care and not as part of a controlled research study or clinical trial.

CIP TCGA Radiology Initiative

Imaging Source Site (ISS) Groups are being populated and governed by participants from institutions that have provided imaging data to the archive for a given cancer type. Modeled after TCGA analysis groups, ISS groups are given the opportunity to publish a marker paper for a given cancer type per the guidelines in the table above. This opportunity will generate increased participation in building these multi-institutional data sets as they become an open community resource. Learn more about the TCGA Breast Phenotype Research Group.

DCE-MRI captures tumour vascular dynamics — pairing it with TCGA molecular data opens up radiogenomic modelling.

Workshop — Breast MRI (DCE-MRI)

Dynamic contrast-enhanced MRI for breast cancer characterization — TCGA-BRCA cohort

Dataset — TCGA-BRCA

Public breast DCE-MRI cohort (TCIA)

COHORT

~108 breast MRI cases in TCGA-BRCA; 84 retained after QC (GE 1.5T, multi-institutional: MSKCC, Mayo, Pittsburgh, Roswell Park; 1999–2004). 88% invasive ductal, 87% ER+, 23% HER2+.

ANALYSIS OPPORTUNITY

Paired imaging + genomic and clinical data (ER/PR/HER2, subtype) enable radiogenomic studies — linking enhancement kinetics and texture to molecular subtype and outcome.

What is DCE-MRI?

Dynamic Contrast-Enhanced MRI

MODALITY

T1-weighted MRI acquired before and repeatedly after IV gadolinium injection. The resulting time series captures contrast wash-in and wash-out kinetics within tissue.

WHY DCE-MRI FOR BREAST CANCER

Tumour neovascularity causes rapid wash-in and washout — highly sensitive for malignant lesions. Standard for high-risk screening, local staging and monitoring response to neoadjuvant therapy.

DCE-MRI captures tumour vascular dynamics — pairing it with TCGA molecular data opens up radiogenomic modelling.

Session 01: Medical Image Formats — DICOM & NIfTI

What is DICOM?

Digital Imaging and Communications in Medicine — the international standard
Every MRI, CT, and X-ray machine in hospitals produces DICOM files.

Patient

└ Study (one clinical visit)

└ Series (one scan type, e.g. T1 post-contrast)

└ Images (individual 2D slices — one .dcm per slice)

Key DICOM Tags

- PatientID, PatientName
- Modality (MR / CT / PT)
- StudyDate, SeriesDescription
- Rows, Columns (image size)
- PixelSpacing (mm/pixel)
- SliceThickness (mm)
- MagneticFieldStrength (T)
- Manufacturer

Images Are Just Numbers

Medical images = NumPy arrays of integers. Any mathematical operation can be applied directly!

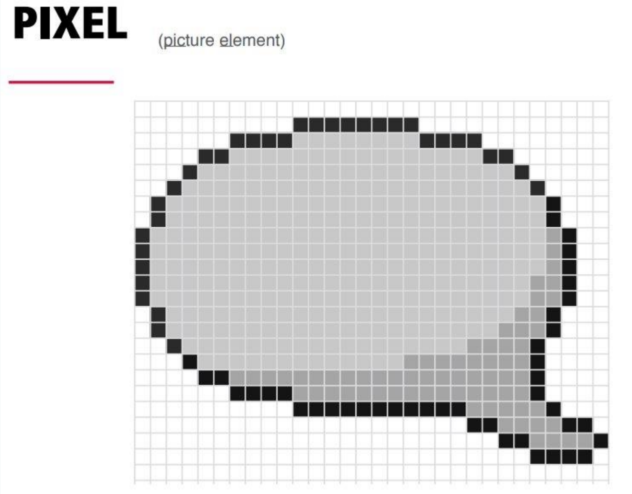

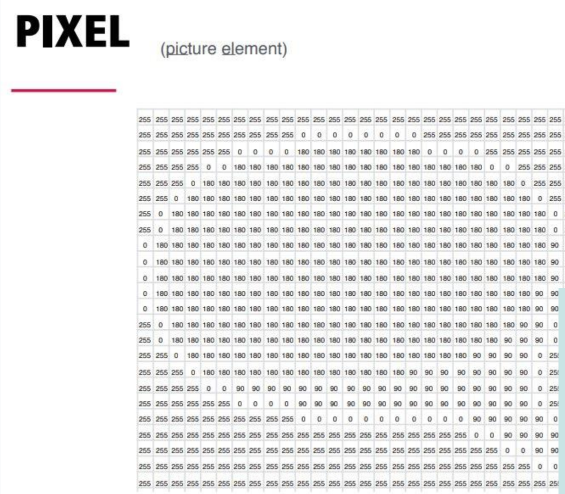

PIXEL (picture element)

255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255
255	255	255	255	255	255	255	255	255	255	0	0	0	0	0	0	0	0	255	255	255	255	255	255	255	255	255	255	255	255
255	255	255	255	255	255	0	0	0	0	180	180	180	180	180	180	180	180	0	0	0	0	255	255	255	255	255	255	255	255
255	255	255	255	0	0	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	0	0	255	255	255	255	255	255
255	255	255	0	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	0	255	255	255	255	255
255	255	0	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	0	255	255	255	255
255	0	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	0	255	255	255
255	0	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	0	255	255	255
0	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	90	0	255	255
0	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	90	0	255	255
0	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	90	0	255	255
0	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	90	90			
0	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	90	90			
255	0	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	90	90	0			
255	0	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	90	90	90	0			
255	255	0	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	180	90	90	90	90	0	255			
255	255	255	0	180	180	180	180	180	180	180	180	180	180	180	180	180	90	90	90	90	90	90	90	90	0	255			
255	255	255	255	0	0	90	90	90	90	90	90	90	90	90	90	90	90	90	90	90	90	90	90	90	0	255			
255	255	255	255	255	255	0	0	0	0	90	90	90	90	90	90	90	90	90	90	90	90	90	90	90	0	255			
255	255	255	255	255	255	255	255	255	255	0	0	0	0	0	0	0	0	0	0	0	90	90	90	90	90	0			
255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	0	0	90	90	90	90			
255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	0	0	90	90			
255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	0	0			
255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255	255			

3D - Voxel

Images Are Just Numbers

Medical images = NumPy arrays of integers. Any mathematical operation can be applied directly!

```
import pydicom
ds = pydicom.dcmread('example.dcm')
image = ds.pixel_array   # NumPy array
print(image.shape)       # (512, 512)
print(image.dtype)       # uint16
print(image.min(), image.max())
```

Transform	Effect
Add 500	Brighter image
Multiply ×2	Higher contrast
Power ^0.5	Square root (gamma)
Log(1+x)	Compress bright values
Clip p1–p99	Remove outliers
Z-score (x-μ)/σ	Standardise scale
Normalize [0,1]	Scale to unit range

NIfTI — The Research Standard

Feature	DICOM	NIfTI
Files per 3D scan	Many (1 per slice)	1 file (.nii.gz)
Used in	Hospital systems	Research & ML
Metadata	Rich clinical tags	Compact header
Python library	**pydicom**	nibabel
Extension	.dcm	.nii / .nii.gz
Convert with	—	dicom2nifti

NIfTI anatomy: Voxel data (3D array) + Affine matrix (voxel→mm) + Header (spacing, dims, dtype)

Visualizing a 3D Volume — Anatomical Planes

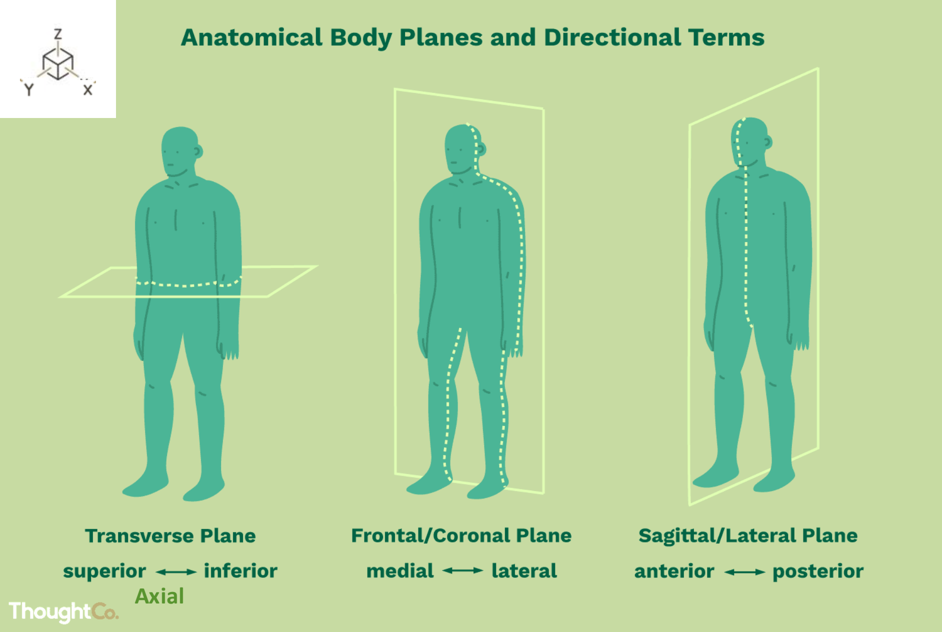

Axial (Z fixed)

- Cross-section from above
- vol[:, :, mid_z]

Coronal (Y fixed)

- View from the front
- vol[:, mid_y, :]

Sagittal (X fixed)

- View from the side
- vol[mid_x, :, :]

```
vol = nib.load('image.nii.gz').get_fdata()   # shape (X, Y, Z)
plt.imshow(vol[:, :, mid_z].T, cmap='gray', origin='lower')   # Axial
plt.imshow(vol[:, mid_y, :].T, cmap='gray', origin='lower')   # Coronal
plt.imshow(vol[mid_x, :, :].T, cmap='gray', origin='lower')   # Sagittal
```

Practical — Session 1 (Notebook 01_medical_image_formats.ipynb)

- Exercise 1 — Explore DICOM metadata
  - Load example.dcm with pydicom
  - Extract: PixelSpacing, SliceThickness, ImageType, MagneticFieldStrength
  - Hint: ds.get('TagName', 'Not available')
- Exercise 2 — Transformation collage
  - Apply 3 different transforms (log, clip, z-score) to the pixel array
  - Display original + 3 transforms in a 1×4 grid with plt.subplots(1,4)
- Exercise 3 — NIfTI exploration
  - Load example.nii.gz with nibabel
  - Print shape, voxel spacing, total physical size in mm
  - Find the 'best slice' in each plane (most non-zero voxels)
  - Plot best axial, coronal, and sagittal slices

Workshop — Breast MRI (DCE-MRI)

Dynamic contrast-enhanced MRI for breast cancer characterization — TCGA-BRCA cohort

Dataset — TCGA-BRCA

Public breast DCE-MRI cohort (TCIA)

COHORT

~108 breast
multi-institu
2004). 88%

ANALYSIS O

Paired imag
enable radi
texture to r

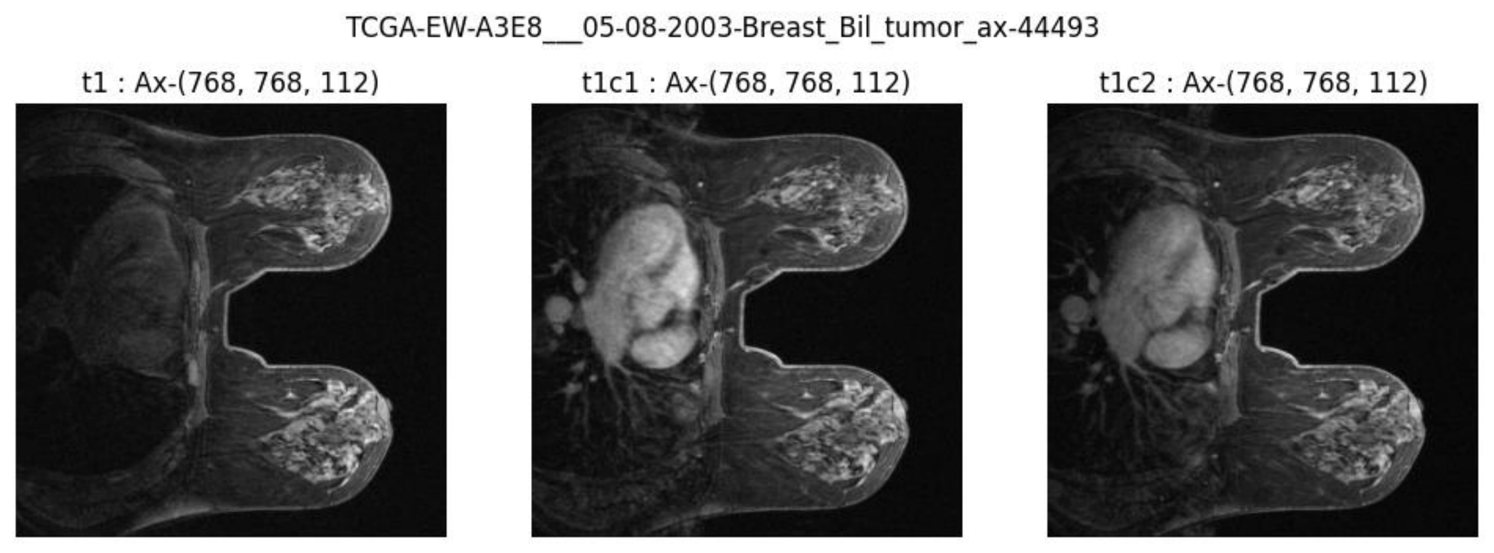

What is DCE-MRI?

Dynamic Contrast-Enhanced MRI

adolinium
n and

— highly
ning, local

DCE-MRI captures tumour vascular dynamics — pairing it with TCGA molecular data opens up radiogenomic modelling.

Workshop — Breast MRI (DCE-MRI)

Dynamic contrast-enhanced MRI for breast cancer characterization — TCGA-BRCA cohort

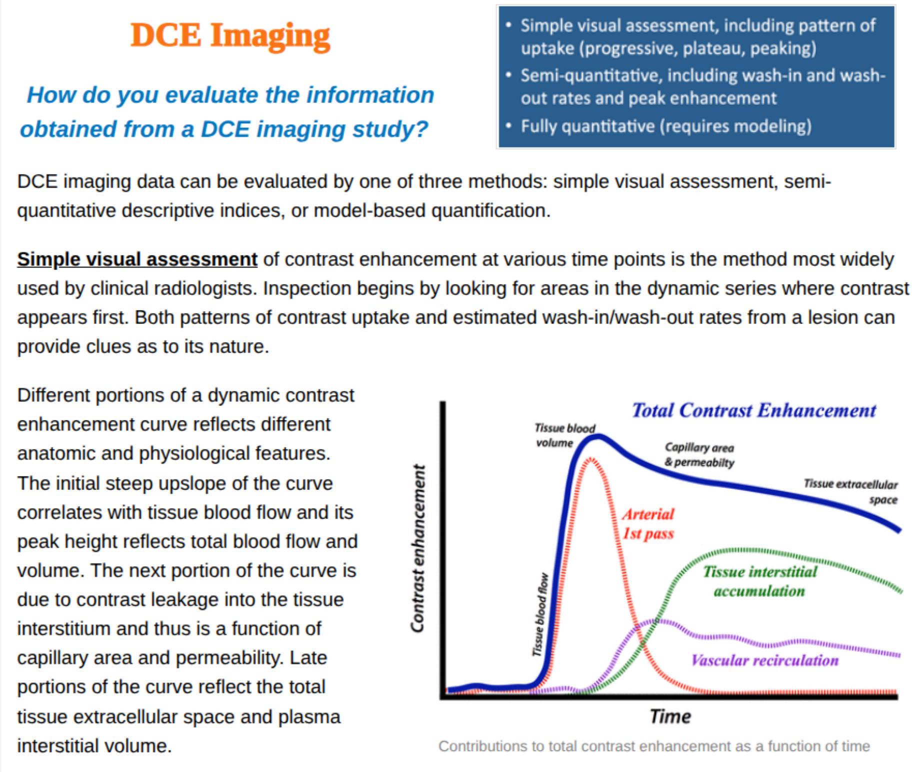

DCE Imaging

How do you evaluate the information obtained from a DCE imaging study?

- Simple visual assessment, including pattern of uptake (progressive, plateau, peaking)
- Semi-quantitative, including wash-in and wash-out rates and peak enhancement
- Fully quantitative (requires modeling)

DCE imaging data can be evaluated by one of three methods: simple visual assessment, semi-quantitative descriptive indices, or model-based quantification.

Simple visual assessment of contrast enhancement at various time points is the method most widely used by clinical radiologists. Inspection begins by looking for areas in the dynamic series where contrast appears first. Both patterns of contrast uptake and estimated wash-in/wash-out rates from a lesion can provide clues as to its nature.

Different portions of a dynamic contrast enhancement curve reflects different anatomic and physiological features. The initial steep upslope of the curve correlates with tissue blood flow and its peak height reflects total blood flow and volume. The next portion of the curve is due to contrast leakage into the tissue interstitium and thus is a function of capillary area and permeability. Late portions of the curve reflect the total tissue extracellular space and plasma interstitial volume.

Contributions to total contrast enhancement as a function of time

Extracted from:
https://mriquestions.com/how-is-dce-analyzed.html

DCE-MRI captures tumour vascular dynamics — pairing it with TCGA molecular data opens up radiogenomic modelling.

Session 02: MRI Preprocessing

Why Preprocess MRI?

- *Raw MRI images contain artefacts that hurt downstream analysis and machine learning.*
- ***Raw MRI data undergo preprocessing to ensure intra- and inter-subject homogeneity before learning tasks.***

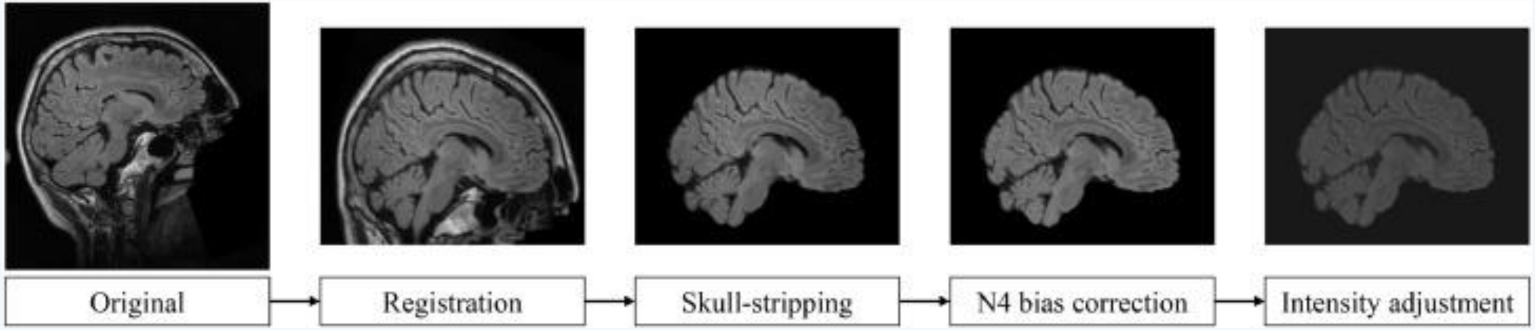

Extracted from 10.1016/j.patrec.2026.02.030

Why Preprocess MRI?

Raw MRI images contain artefacts that hurt downstream analysis and machine learning.

Problem	Cause	Fix (this session)
Intensity inhomogeneity	Magnetic field non-uniformity	N4 Bias Correction
Noise / speckle	Thermal noise in receiver coils	NLMeans Denoising
Intensity outliers	Bright arteries, metal implants	Percentile Clipping
Different intensity scales	Different scanners / protocols	Z-score Normalisation ⚠
Different resolutions	Different scanners	Resampling to 1×1×1 mm

⚠ Z-score normalisation is NOT recommended for DCE-MRI — it destroys contrast dynamics!

Step — N4 Bias Field Correction

MRI scanners produce smooth intensity variations (bias field) across the field of view.
N4 estimates and removes this background variation so the same tissue looks uniform.

How it works

The algorithm estimates the intensity distribution of the image and applies a correction factor to reduce the intensity inhomogeneities. The correction factor is based on the intensity histogram of the image, and it is applied using a convolution operation.

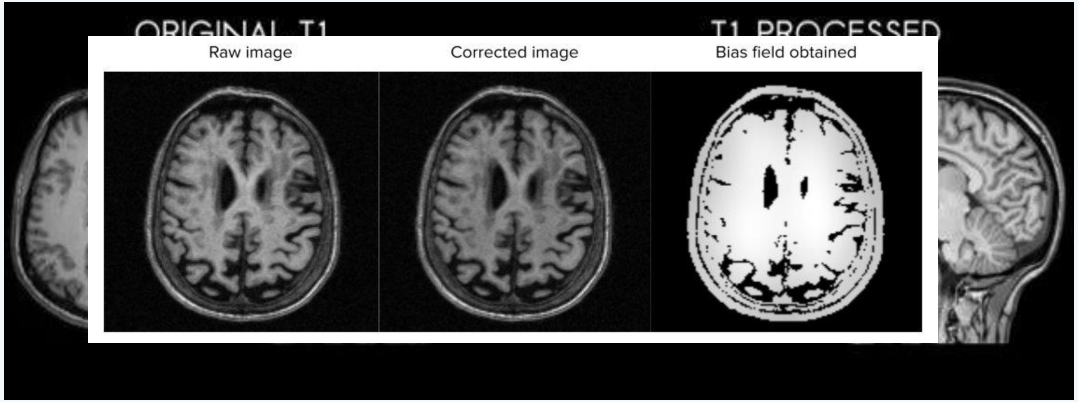

Extracted from https://younsess-elbrag.medium.com/post-processing-medical-images-n4-bias-correction-and-skull-stripping-572a10b4b3

```
corrected = bias_correction(image_sitk, shrink_factor=4)
# Uses: sitk.N4BiasFieldCorrectionImageFilter  →  image / exp(log_bias_field)
```

What is an image histogram?

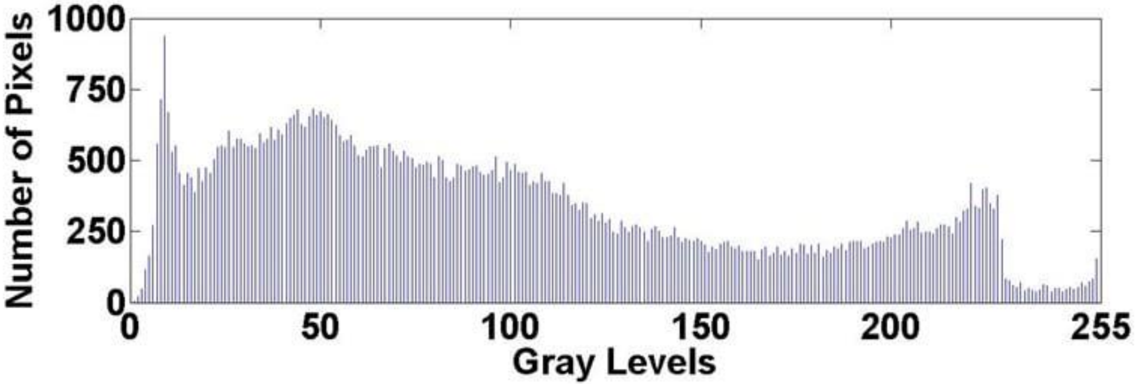

(a)

(b)

Step — Clipping & Intensity Normalization

Percentile Intensity Clipping

- Removes extreme outlier voxels (bright vessels, metal)
- Compute 1st and 99th percentile of non-zero voxels
- Clamp: values < p1 → p1, values > p99 → p99
- ✓ Safe for DCE-MRI (only removes outliers)
- p5–p95 → aggressive, loses some info
- p1–p99 → conservative (recommended)

```
clipped  = clip_intensity(denoised, lower_pct=1, upper_pct=99)
# Then optionally: normalised = zscore_norm(clipped)  ← skip for DCE-MRI!
```

Steps — Clipping & Intensity Normalization

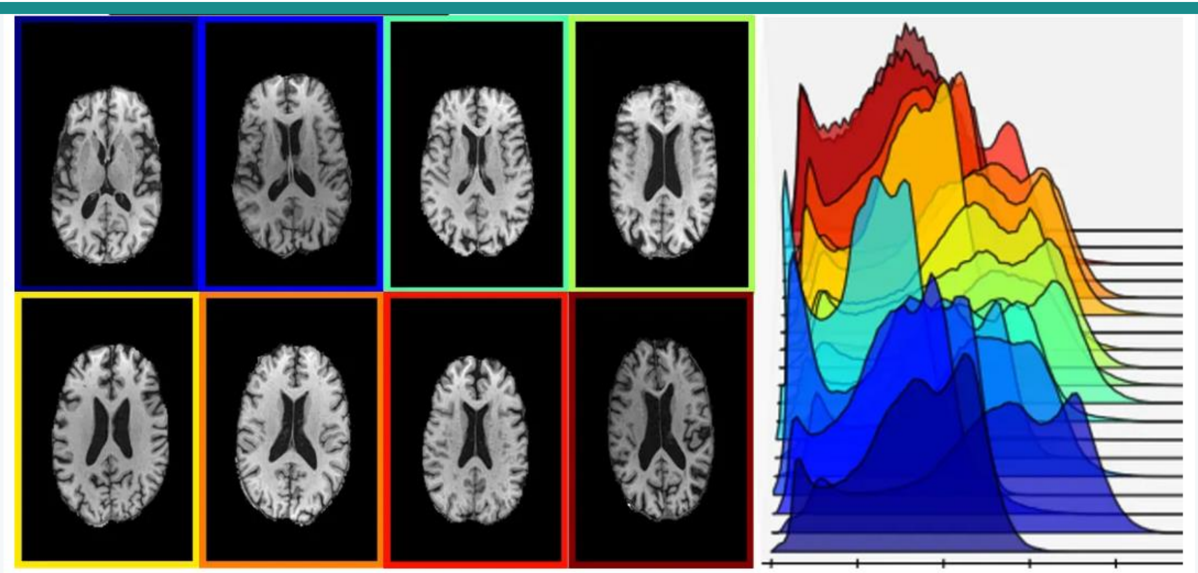

Conventional MRI has exceptional soft tissue contrast and it allows the radiologist to differentiate between different tissue types.

Extracted from https://medium.com/@susanne.schmid/image-normalization-in-medical-imaging-f586c8526bd1

```
clipped  = clip_intensity(denoised, lower_pct=1, upper_pct=99)
# Then optionally: normalised = zscore_norm(clipped)  ← skip for DCE-MRI!
```

Steps — Clipping & Intensity Normalization

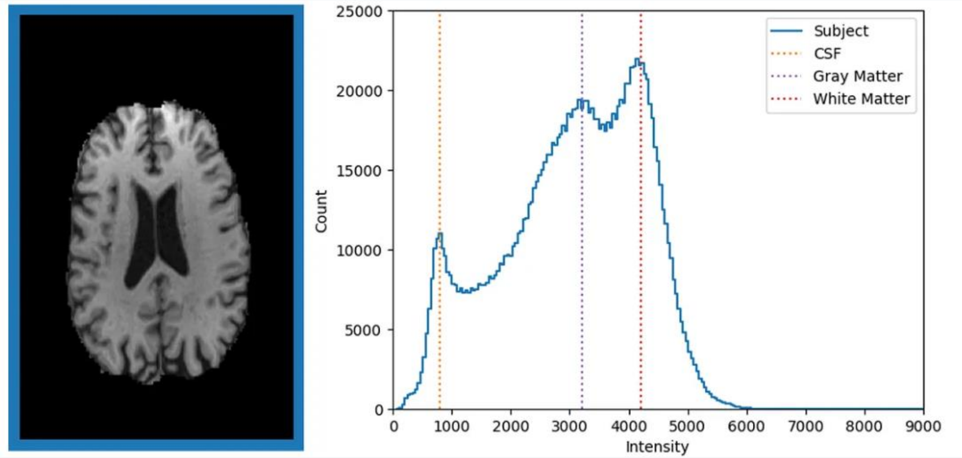

Cerebrospinal fluid (CSF)

Looking at the histograms of the entire volume, allows us to see that certain tissues are more noticeable than others. This is different compared to natural images where the intensity value reflects the brightness of the object.

Extracted from https://medium.com/@susanne.schmid/image-normalization-in-medical-imaging-f586c8526bd1

```
clipped  = clip_intensity(denoised, lower_pct=1, upper_pct=99)
# Then optionally: normalised = zscore_norm(clipped)  ← skip for DCE-MRI!
```

Steps — Clipping & Intensity Normalization

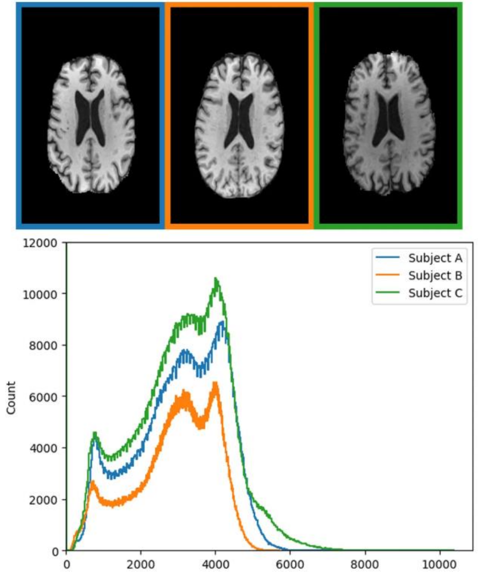

Conventional MRI is qualitative which means, that even though the intensity value refers to a specific tissue, a repeated MRI scan will obtain different voxel intensities. However, the absolute intensity does not matter in MRI because MR images benefit from the contrast between the different tissues and not the absolute value.

On the left we have figure of three scans and their histograms acquired with the same T1 pulse sequence at a GE scanner.

Extracted from https://medium.com/@susanne.schmid/image-normalization-in-medical-imaging-f586c8526bd1

```
clipped  = clip_intensity(denoised, lower_pct=1, upper_pct=99)
# Then optionally: normalised = zscore_norm(clipped)  ← skip for DCE-MRI!
```

Steps — Clipping & Intensity Normalization

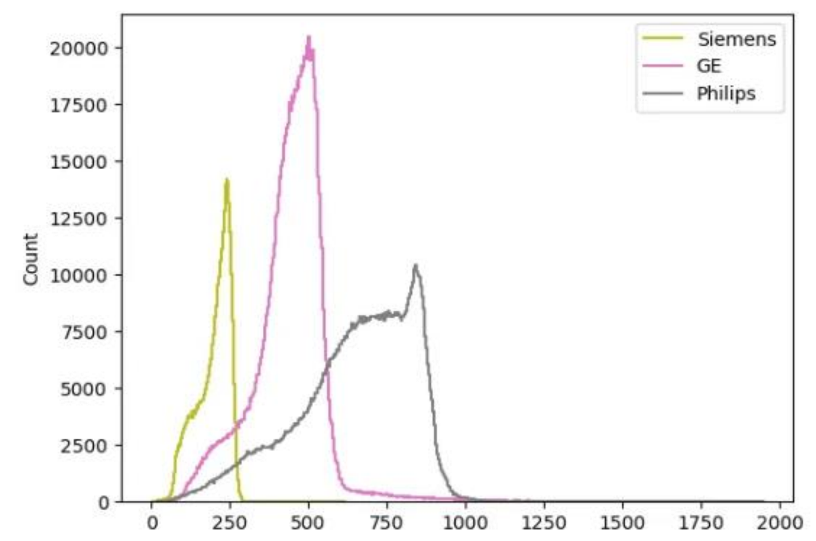

Some scanners have images with an absolute intensity value of 10,000 others only have intensities of 6,000. In the figure below, histograms of T1 images for different vendors are visualized.

Extracted from https://medium.com/@susanne.schmid/image-normalization-in-medical-imaging-f586c8526bd1

```
clipped  = clip_intensity(denoised, lower_pct=1, upper_pct=99)
# Then optionally: normalised = zscore_norm(clipped)  ← skip for DCE-MRI!
```

Why Clip & Normalise MRI Intensities?

1. Why normalise at all — it's about optimisation, not aesthetics

- Inputs on different scales → unstable gradients (vanishing or exploding)
- Loss landscape becomes harder to optimise; training is slower and less stable
- Normalisation reduces voxel-value variance, smooths convergence, stabilises updates

2. MRI breaks the assumptions behind Min-Max

- Intensity reflects tissue type, not brightness — WM, GM, CSF, fat each form histogram peaks
- Intensity is relative — same scanner, same patient, repeat scan → different absolute values
- Intensity is unbounded — no fixed ceiling; ranges shift from ~5 000 to 10 000+ across scans

3. Why clipping comes BEFORE normalisation

- Outliers (motion spikes, fat artefacts, hyper-intense lesions) dominate min/max → unstable scaling
- Clipping at the 1st–99th percentile removes tail noise while preserving tissue contrast
- Result: normalisation operates on the meaningful signal range, not on artefacts

Can we find some ideas for normalization?

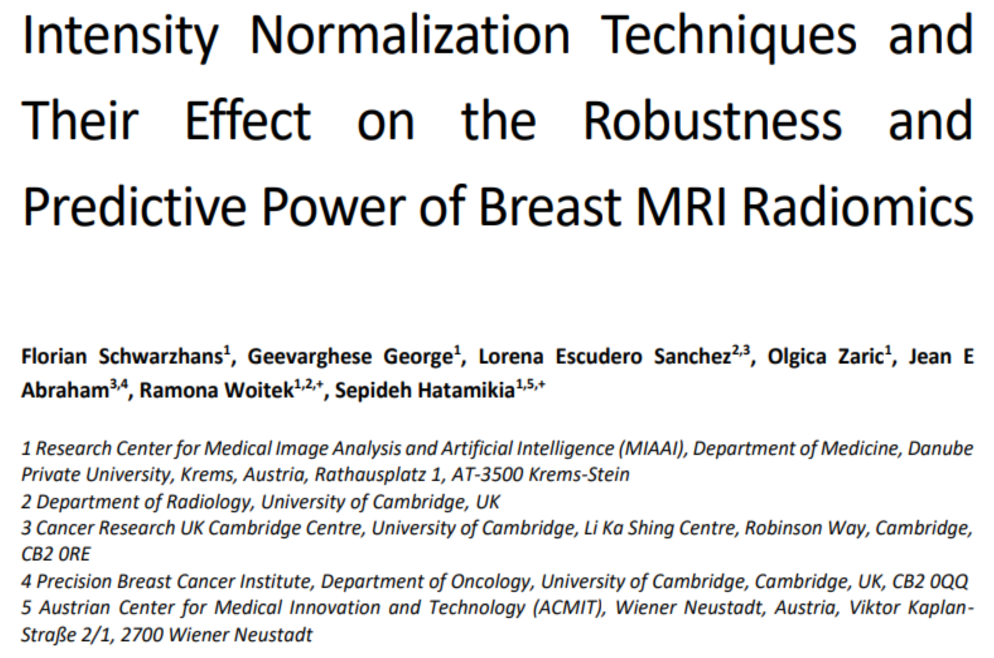

Intensity Normalization Techniques and Their Effect on the Robustness and Predictive Power of Breast MRI Radiomics

Florian Schwarzhans[1], Geevarghese George[1], Lorena Escudero Sanchez[2,3], Olgica Zaric[1], Jean E Abraham[3,4], Ramona Woitek[1,2,+], Sepideh Hatamikia[1,5,+]

1 Research Center for Medical Image Analysis and Artificial Intelligence (MIAAI), Department of Medicine, Danube Private University, Krems, Austria, Rathausplatz 1, AT-3500 Krems-Stein
2 Department of Radiology, University of Cambridge, UK
3 Cancer Research UK Cambridge Centre, University of Cambridge, Li Ka Shing Centre, Robinson Way, Cambridge, CB2 0RE
4 Precision Breast Cancer Institute, Department of Oncology, University of Cambridge, Cambridge, UK, CB2 0QQ
5 Austrian Center for Medical Innovation and Technology (ACMIT), Wiener Neustadt, Austria, Viktor Kaplan-Straße 2/1, 2700 Wiener Neustadt

doi.org/10.1016/j.ejrad.2025.112086

Min–Max Normalization

- Rescales to [0, 1]
- 1%–99% variant was **discarded** — saturated lesion values

Z-Score Normalization (3 variants)

- **ZA** — whole image
- **ZB** — breast tissue mask
- **ZL** — lesion only

Piecewise Linear Histogram Equalization (PLHE)

- Non-linear, histogram-based (Nyúl method)
- Reference histogram built from a random subset of scans

Can we find some ideas for normalization?

The paper uses only the **first post-contrast phase** — treating DCE volumes as **static T1-weighted post-contrast images**. The dynamic (temporal) nature of the sequence is not exploited.

⚠ Risks of aggressive normalization

DCE intensities encode **physiologically meaningful** wash-in / wash-out kinetics. Aggressive normalization can:

- Destroy **relative enhancement** between timepoints
- Break the link to **contrast agent dynamics**
- Prevent within-patient comparison of enhancement curves

✓ What to preserve

For DCE to stay interpretable per-patient:

- Normalize against a **pre-contrast baseline** (e.g. post / pre signal ratio) — scanner-invariant
- Keep **relative** (not absolute) intensities across timepoints
- Use a **reference tissue** (e.g. muscle, fat) for cross-patient scaling

Why Coregistration Matters

Multi-parametric MRI features only make sense when T1, T2 and DCE describe the same voxel.

Coregistration finds a **spatial transformation** that aligns all timepoints to a common reference (usually t0, pre-contrast).

1. Rigid registration
- Only allows **translation + rotation** (6 degrees of freedom)
- Fast and simple
- Good for: brain MRI (skull constrains motion)

2. Affine registration
- Translation + rotation + **scaling + shearing** (12 degrees of freedom)
- Still linear
- Better than rigid but still limited for soft tissue

3. Deformable (non-rigid) registration ← recommended for breast
- Allows **local deformations** of the image
- Models actual tissue deformation from breathing etc.
- Much more accurate for breast DCE-MRI
- Common algorithms: **B-spline**, **Demons**, **SyN (ANTs)**

https://andysbrainbook.readthedocs.io/en/latest/SPM/SPM_Short_Course/SPM_04_Preprocessing/03_SPM_Coregistration.html

Step — Isotropic Resampling

Different scanners produce different voxel sizes. Resampling to 1×1×1 mm makes images comparable and ensures radiomic features are resolution-independent.

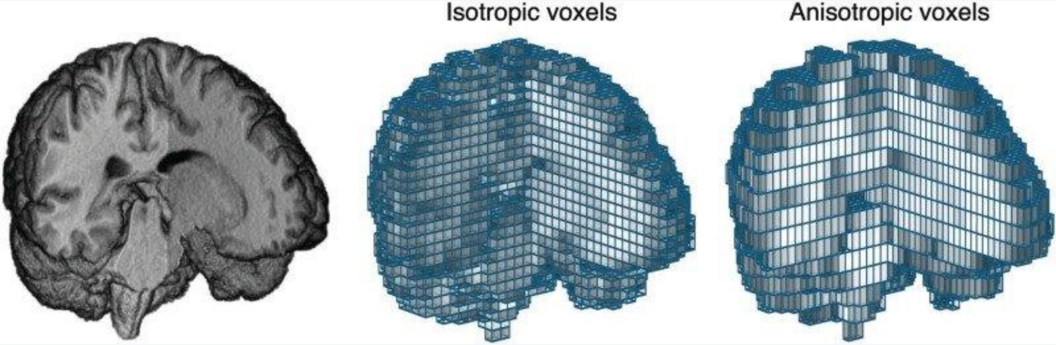

	After Resampling
0.7	1.0 × 1.0 × 1.0 mm (isotropic)
Di	Comparable feature values ✓

```
resampled = resample(image_sitk, new_spacing=(1.0, 1.0, 1.0))
# Computes new_size = original_size × original_spacing / new_spacing
# Uses BSpline interpolation for image quality
# Recommended interpolator: sitk.sitkBSpline (image), sitk.sitkNearestNeighbor (mask)
```

Recommended Pipeline for DCE-MRI

Raw MRI

↓

N4 Bias Correction ✓ always

↓

NLMeans Denoising ✓ optional

↓

Intensity Clipping ✓ safe for DCE

↓

Z-score Norm ⚠ skip for DCE?!

↓

Coregisteration and Resample to 1×1×1 mm ✓ for radiomics

Practical — Session 2 (Notebook 02_mri_preprocessing.ipynb)

- Exercise 1 — Clip with different percentiles
  - Apply clip_intensity with: (5,95), (10,90), (0.1,99.9)
  - Show all 4 results side by side (original + 3 clips)
  - Which is most aggressive? Which is most conservative?
- Exercise 2 — Full preprocessing pipeline
  - Load example_t1pre.nii.gz
  - Chain: bias_correction → nlmeans_denoise → clip_intensity → resample(1,1,1)
  - Use show_before_after(raw, step4, 'Full pipeline') to compare
  - Skip Z-score normalisation for DCE-MRI!
- Key parameters to experiment with:
  - shrink_factor: 2 vs 4 vs 8 (speed vs accuracy of bias correction)
  - h: 0.5 vs 0.8 vs 1.2 (denoising strength)
  - new_spacing: (1,1,1) vs (0.5,0.5,1) (resolution tradeoffs)

Session 03: Radiomics Feature Extraction

What is Radiomics?

Extracting **quantitative features** from medical images to capture tissue and lesion properties — shape, heterogeneity, texture — for **clinical decision support**.

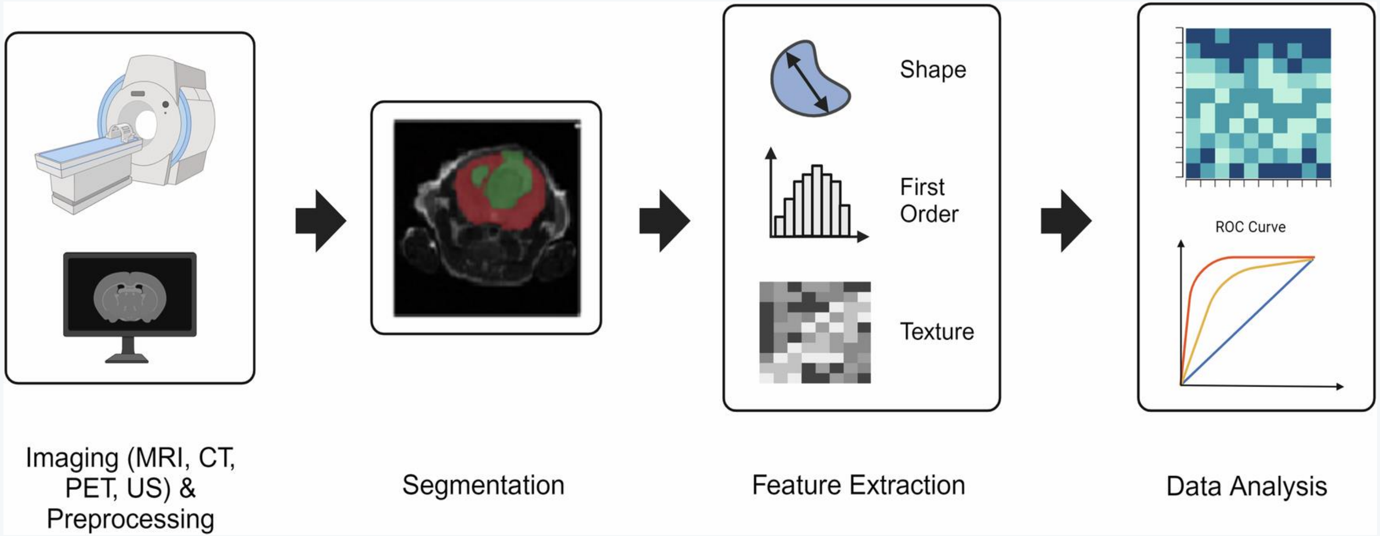

Sources: doi.org/10.2967/jnumed.118.222893 · doi.org/10.1038/s44303-025-00104-z

What is Radiomics?

Radiomics = extraction of large numbers of quantitative features from medical images.
Instead of visual inspection, radiomics converts images to numbers that describe shape, intensity, and texture — ready for machine learning.

Image + Mask → Preprocessing → **Feature Extraction (PyRadiomics)** → Feature Table → Machine Learning

Feature categories extracted by PyRadiomics

- Shape (~14): **segmentation/lesion** -> volume, sphericity, compactness, surface area, elongation...
- First-order (~18): **describe pixel intensity** -> mean, std, skewness, kurtosis, entropy, energy...

Texture feature families:

- GLCM (~24): co-occurrence texture patterns between voxel pairs
- GLRLM (~16): runs of equal grey-levels in each direction
- GLSZM (~16): zones of equal grey-levels
- GLDM (~14): grey-level dependence matrix
- NGTDM (~5): neighbourhood tone differences → Total: ~100–200 features per image!

Backup - Radiomics Texture Families

Feature categories extracted by PyRadiomics (Cont.)

Texture feature families:

- GLCM (~24): co-occurrence texture patterns between voxel pairs
- GLRLM (~16): runs of equal grey-levels in each direction
- GLSZM (~16): zones of equal grey-levels
- GLDM (~14): grey-level dependence matrix
- NGTDM (~5): neighbourhood tone differences → Total: ~100–200 features per image!

“How often does a dark voxel appear next to a bright voxel?”
Measures consecutive runs of voxels with the same intensity.

- Measures connected zones of equal intensity, independent of direction (Quantifies clustered regions with similar intensity)

- How strongly surrounding voxels resemble the center voxel. (Measures local intensity similarity.)

Feature Extraction with PyRadiomics

```python
from radiomics import featureextractor

extractor = featureextractor.RadiomicsFeatureExtractor()
extractor.enableAllFeatures()

result = extractor.execute('image.nii.gz', 'mask.nii.gz')

# Filter to numeric features only (remove diagnostics)
features = {k: float(v) for k, v in result.items()
            if not k.startswith('diagnostics') and not isinstance(v, str)}

# Example feature names:
# original_shape_VoxelVolume, original_firstorder_Mean, original_glcm_Contrast ...
```

Typical output: 100–200 features · Takes 30–60 s per image · Features named as: original_{category}_{name}

Practical — Session 3 (Notebook 03_radiomics_features.ipynb)

- Exercise 1 — Enhancement metrics (inside tumour ROI)
  - Compute mean signal per phase: mean_pre, mean_c1, mean_c2
  - Absolute enhancement at c1: mean_c1 − mean_pre
  - Relative enhancement at c1: (mean_c1 − mean_pre) / mean_pre × 100 [%]
  - Washout between phases: mean_c2 − mean_c1
  - Hint: arr[mask_arr].mean() gives voxels inside the ROI
- Exercise 2 — Features from different phases
  - Extract features from t1pre.nii.gz AND t1c1.nii.gz with the same mask
  - Compute absolute difference for each common feature
  - Which feature category shows the largest change between phases?
- Exercise 3 (Bonus) — Feature correlation heatmap
  - Build a DataFrame with pre and c1 features as rows
  - Plot seaborn.heatmap of first-order feature correlations across phases

Session 04: Predictive Modelling with Radiomics

Dataset — TCGA-BRCA Breast Cancer Radiomics

Goal: predict pathological cancer stage from MRI radiomic features.
Data: 91 patients from The Cancer Genome Atlas Breast Cancer (TCGA-BRCA) cohort.

Radiomics file (XLS)

- 91 patients × multiple features
- Pre-computed workstation features
- Source: UChicago V2010 MRI workstation
- Reference: TCGA Breast Radiogenomics 2014
- URL: cancerimagingarchive.net

Clinical file (TSV)

- age_at_initial_pathologic_diagnosis
- er_status_by_ihc (ER receptor)
- pr_status_by_ihc (PR receptor)
- her2_status_by_ihc_in_percent_positive
- pathologic_stage → grouped to Stage I/II/III

Stage grouping: Stage IA/IB → Stage I · Stage IIA/IIB → Stage II · Stage IIIA/IIIC → Stage III

Machine Learning Workflow

Merge data
Radiomics + Clinical → Impute missing
(median) → Scale features
(StandardScaler) → LOOCV
Leave-One-Out CV → Evaluate
Accuracy

Why Leave-One-Out Cross-Validation (LOOCV)?

- Only 91 patients → standard train/test split wastes too much data
- LOOCV trains on N−1 patients, tests on 1, repeats N times
- Every patient is used for testing exactly once
- Gives a robust accuracy estimate for small datasets
- Sklearn: `cross_val_score(pipe, X, y, cv=LeaveOneOut(), scoring='accuracy')`

Three Models — Comparing Feature Sources

Radiomics only

- Imaging features
- extracted from MRI

Clinical only

- Age, ER status,
- PR status, HER2 status

Radiomics + Clinical (combined)

- All radiomic features
- + all clinical features
- (np.hstack)

```
pipe = Pipeline([('scaler', StandardScaler()), ('clf',
RandomForestClassifier(n_estimators=100))])
loo  = LeaveOneOut()
y_pred = cross_val_predict(pipe, X, y, cv=loo)
acc = accuracy_score(y, y_pred)
bal_acc = balanced_accuracy_score(y, y_pred)
f1_mac = f1_score(y, y_pred, average="macro", zero_division=0)
```

Exploratory Data Analysis — Always First!

- Step 1 — Clinical overview
  - Bar chart: cancer stage distribution (Stage I / II / III)
  - Histogram: age at diagnosis
  - Pie chart: ER receptor status (Positive / Negative / Indeterminate)
- Step 2 — Features by stage
  - Violin plots: feature distributions grouped by stage_group
  - Look for features that differ across stages → potential predictors
  - Visualise first 12 features with sns.violinplot(data=df, x='stage_group', y=feat)
- Step 3 — Correlation heatmap
  - feat_data.corr() → pairwise feature correlations
  - sns.heatmap(corr, cmap='RdBu_r', center=0) → red = positive, blue = negative
  - Many radiomic features are highly correlated → feature selection is important!

Practical — Session 4 (Notebook 04_predictive_modeling.ipynb)

- Exercise 1 — Interpret model coefficients
  - Fit the model with all data
  - Get importances from the trained model
  - Anything unexpected?

  Exercise 2 — Change model hyperparameters
  - Include a hyperparameter related to the imbalance class in the make_pipeline() function

  Exercise 3 — Try Logistic Regression
  - Replace RandomForestClassifier with LogisticRegression(max_iter=1000)
  - Evaluate with LOOCV on the radiomics feature matrix
  - Compare accuracy vs Random Forest — which is better?
  - Note: simpler models often win on small datasets!

⚠ Important Caveats for Real Radiomics Research

- Small dataset (n=91): high risk of overfitting — these results are educational, not clinical
- Feature selection is critical: 100+ features with 91 patients → use LASSO, RFE, or univariate tests
- External validation required: train on one cohort, test on an independent cohort
- Reproducibility: PyRadiomics settings (bin width, interpolation) must match across sites
- Clinical expert needed: findings must be interpreted alongside clinical context
- Normalisation consistency: same preprocessing pipeline must be applied at inference time
- Class imbalance: Stage I vs Stage III counts may differ → consider balanced accuracy or SMOTE
- Real datasets need: hundreds to thousands of patients for robust radiomic signatures

⚠ Important Caveats for Real Radiomics Research

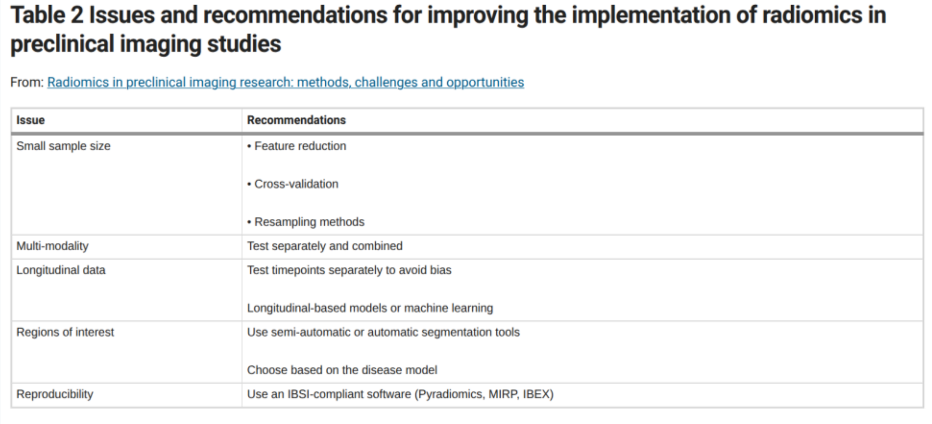

Table 2 Issues and recommendations for improving the implementation of radiomics in preclinical imaging studies

From: Radiomics in preclinical imaging research: methods, challenges and opportunities

Issue	Recommendations
Small sample size	• Feature reduction • Cross-validation • Resampling methods
Multi-modality	Test separately and combined
Longitudinal data	Test timepoints separately to avoid bias Longitudinal-based models or machine learning
Regions of interest	Use semi-automatic or automatic segmentation tools Choose based on the disease model
Reproducibility	Use an IBSI-compliant software (Pyradiomics, MIRP, IBEX)

Why it helps

- Non-invasive readout of disease and treatment over time
- Pairs with ML for detection, classification & response prediction
- Multi-modal & longitudinal — links imaging to histology

Watch out for

- **Small samples** + high-dimensional features → overfitting
- Mitigate: feature reduction, cross-validation, resampling
- Multi-modal / longitudinal data: assess dependence before pooling

⚠ Important Caveats for Real Radiomics Research

- Small dataset (n=91): high risk of overfitting — these results are educational, not clinical
- Feature selection is critical: 100+ features with 91 patients → use LASSO, RFE, or univariate tests
- External validation required: train on one cohort, test on an independent cohort
- Reproducibility: PyRadiomics settings (bin width, interpolation) must match across sites
- Clinical expert needed: findings must be interpreted alongside clinical context
- Normalisation consistency: same preprocessing pipeline must be applied at inference time
- Class imbalance: Stage I vs Stage III counts may differ → consider balanced accuracy or SMOTE
- Real datasets need: hundreds to thousands of patients for robust radiomic signatures

Extra slides on feature extraction

Project 2 — Brain MRI · Part 2: Proposed Solution

Feature-extraction + machine learning pipeline for surgical decision support and prognosis prediction

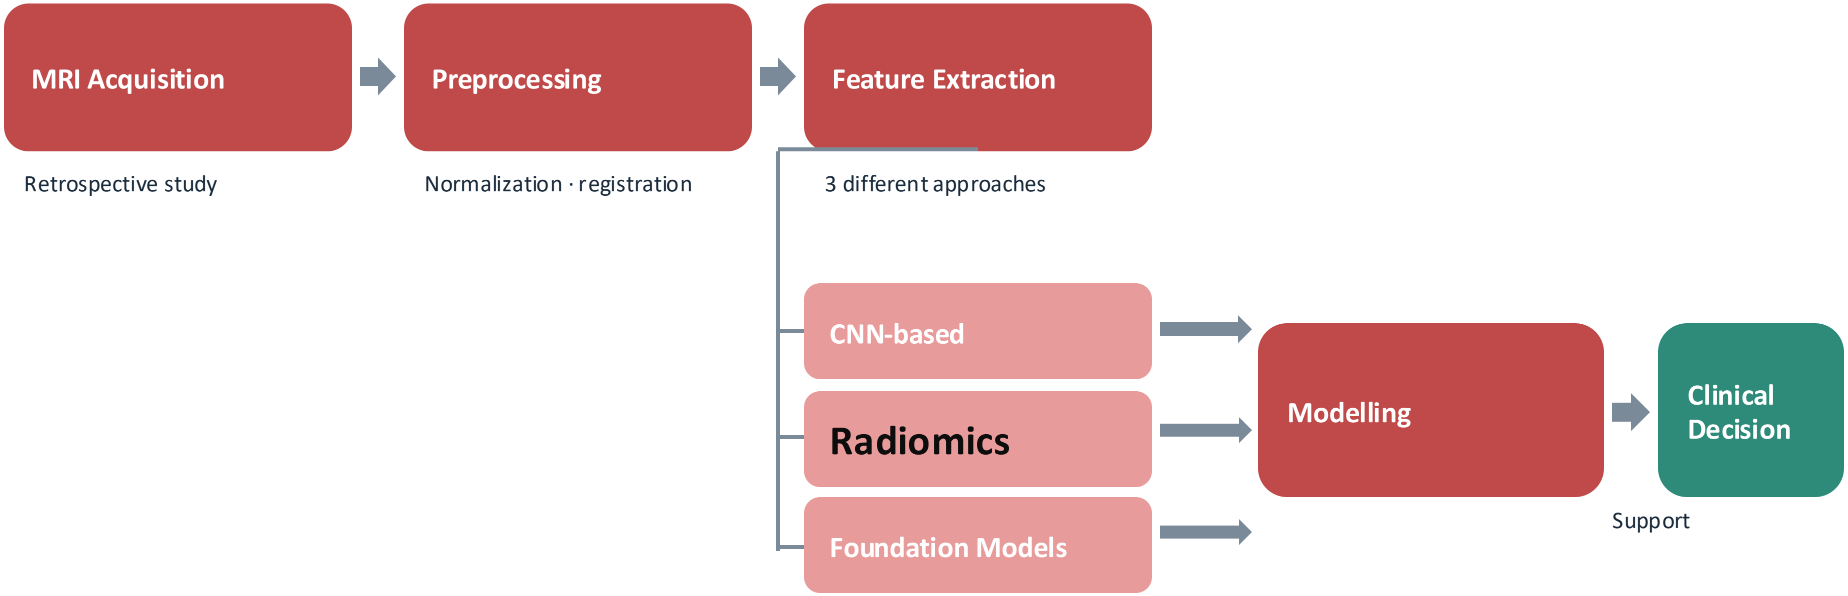

MRI Acquisition

Retrospective study

Preprocessing

Normalization · registration

Feature Extraction

3 different approaches

CNN-based

Radiomics

Foundation Models

Modelling

Clinical Decision

Support

End-to-end pipeline — from raw MRI to support clinical decision

Extra Slides: Feature Selection

In this workshop we have mainly considered a classification problem (through supervised learning) where features were handcrafted. We told in our pipeline that the features were X, Y, and Z ….

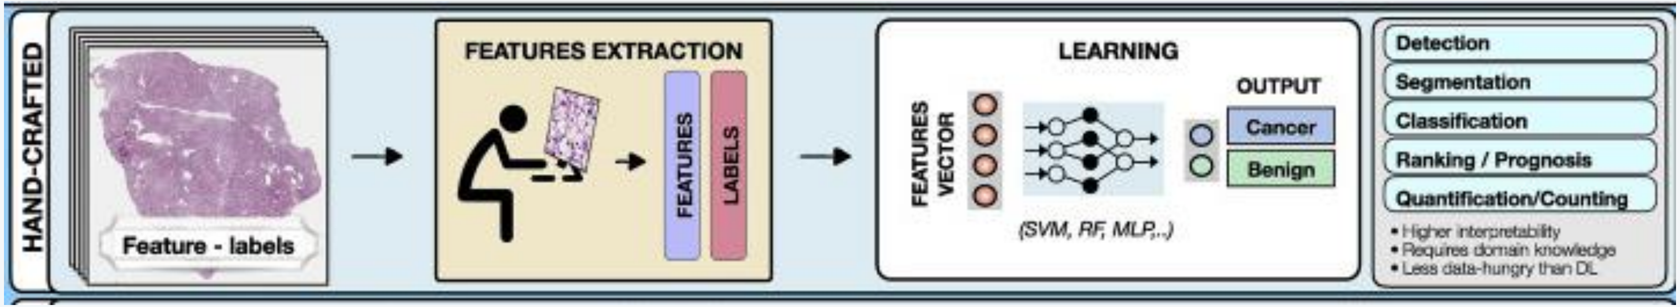

Could the machine / computers connect and understand the image in a different way? → YES

Lipkova J et al., Artificial intelligence for multimodal data integration in oncology, Cancer Cell, 40, 1095-1110

Extra Slides: Feature Selection

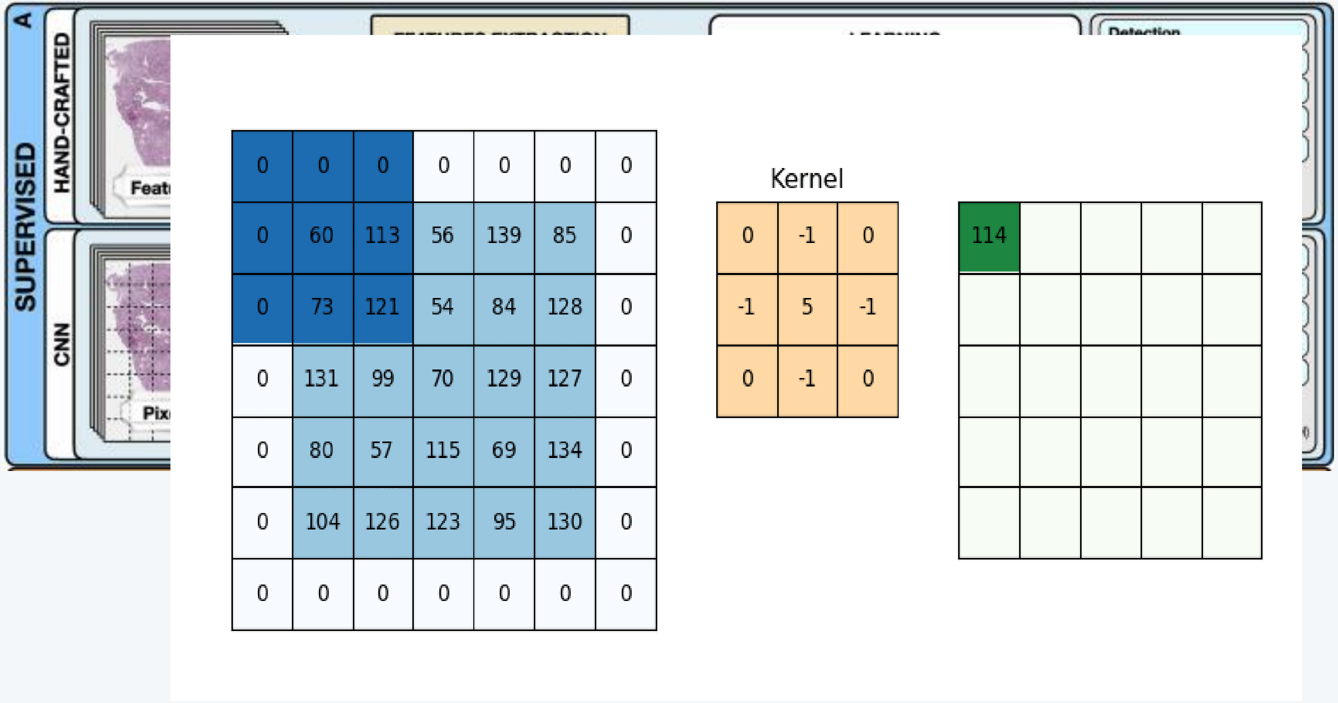

0	0	0	0	0	0	0
0	60	113	56	139	85	0
0	73	121	54	84	128	0
0	131	99	70	129	127	0
0	80	57	115	69	134	0
0	104	126	123	95	130	0
0	0	0	0	0	0	0

Kernel

0	-1	0
-1	5	-1
0	-1	0

114				

From https://medium.com/@Suraj_Yadav/in-depth-knowledge-of-convolutional-neural-networks-b4bfff8145ab

Lipkova J et al., Artificial intelligence for multimodal data integration in oncology, Cancer Cell, 40, 1095-1110

Extra Slides: Feature Selection

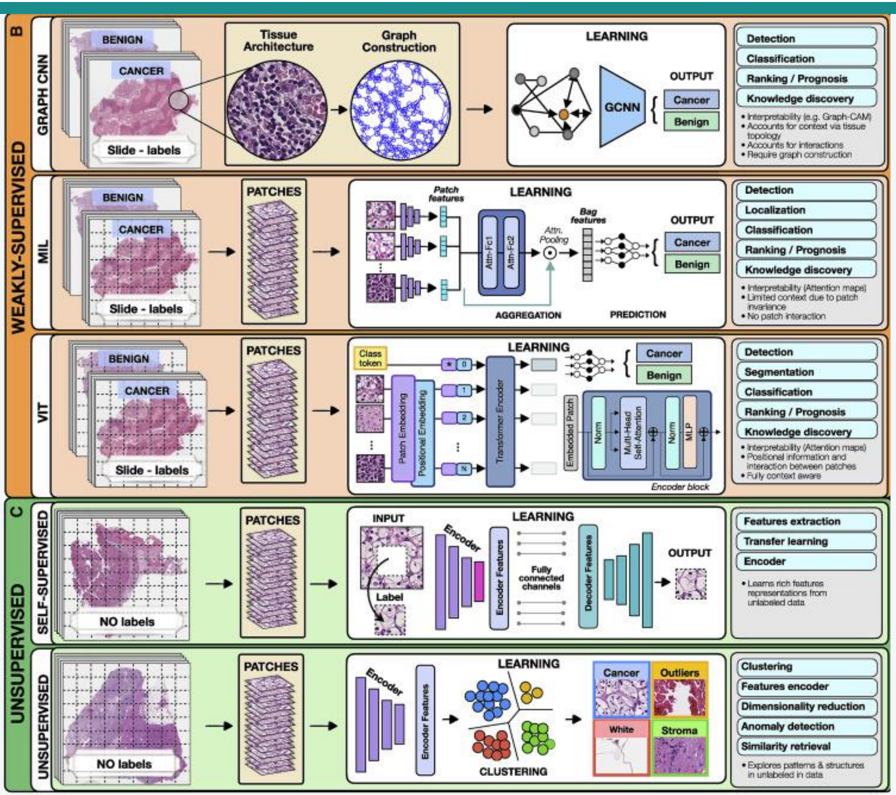

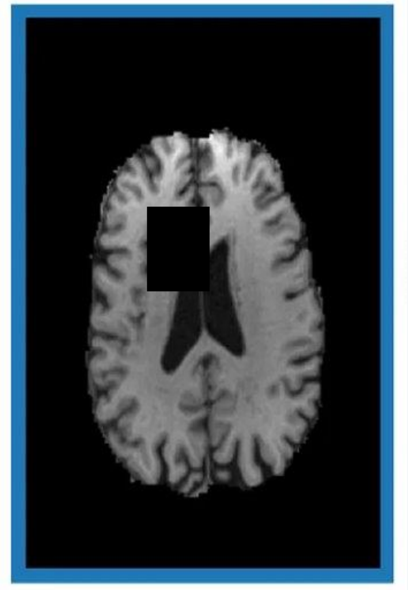

Lipkova J et al., Artificial intelligence for multimodal data integration in oncology, Cancer Cell, 40, 1095-1110

Workshop Complete!

S01 Medical Image Formats	DICOM = hospital standard (1 slice/file); NIfTI = research (1 file for full 3D volume)
S02 MRI Preprocessing	N4 bias → NLMeans denoise → clip → resample.
S03 Radiomics Features	PyRadiomics extracts 100–200 features (shape, first-order, texture) per image
S04 Predictive Modelling	LOOCV on small datasets; combined radiomics + clinical often wins

Thank you! Questions welcome.