Medical Image Processing

Guide to AI and Data Courses · Part III: Signals, Images and Audio

Imaging modalities | DICOM and 3D volumes | Segmentation | Registration | Diagnostic support | Related courses

Sercan KÜLCÜ · Giresun University · Computer Engineering

© Sercan Külcü. All rights reserved. · Interactive HTML version of these slides:
https://sercankulcu.github.io/files/ai_courses/slides/Chapter_20_Medical_Image_Processing.html

INTRODUCTION

What Question Does It Ask?

> **How do X-ray, CT and MRI images help diagnosis and treatment?**

Medical image processing enhances, segments, measures and compares X-ray, CT, MRI and ultrasound images to support **diagnosis and treatment**. It applies image processing and computer vision methods under the constraints of little data, expensive labels and patient safety.

At a glance

APPROACH image processing plus clinical knowledge	**INPUT** CT, MRI, X-ray, ultrasound	**OUTPUT** segmentation, measurement, diagnostic suggestion
SUCCESS MEASURE sensitivity, specificity, Dice	**MATH LOAD** medium-high	**DATA NEEDS** little; labels are expensive

A Short History

1895
Röntgen discovers X-rays; the first medical image is an X-ray of his wife's hand.

1972
Hounsfield's computed tomography (CT) scanner is introduced: cross-sections are obtained by computation.

1973
Lauterbur publishes the first imaging paper using magnetic resonance, laying the foundation of MRI.

1993
The DICOM 3.0 standard is published; devices from different vendors speak the same file format.

2015
U-Net is published: an encoder-decoder network for segmentation with few labeled medical images.

2018
The US FDA authorizes IDx-DR, which screens for diabetic retinopathy without a physician.

THE COURSE

Core Concepts

Imaging modalities
X-ray, CT, MRI, ultrasound, PET: each a different physics.

DICOM
A standard that carries the image and patient data together.

Voxels and volumes
CT and MRI are 3D; voxels instead of pixels.

Denoising
Preserving quality in low-dose CT and in MRI.

Segmentation
Outlining an organ, tumor or vessel.

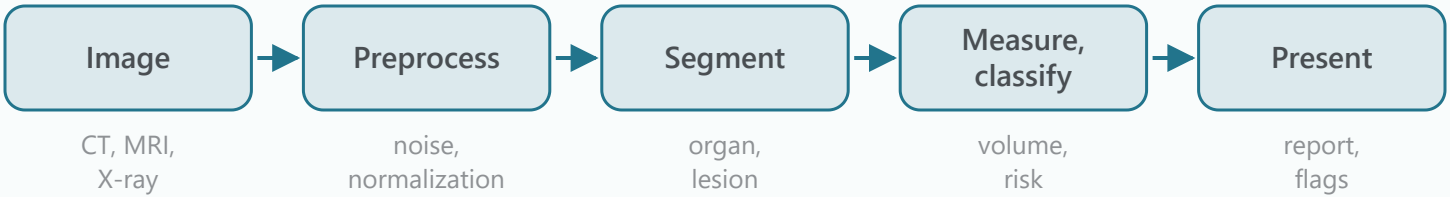
Registration
Aligning two images (before/after, CT/MRI).

Aided diagnosis (CAD)
The computer flags suspicious regions for the physician.

Clinical evaluation
Sensitivity, specificity, Dice; comparison with physicians.

THE COURSE

A Typical Workflow



Example: Hundreds of slices from a lung CT scan are read as DICOM, the lung region is segmented and a network flags suspicious nodules. Registered against the previous scan, the growth in nodule volume is measured. The radiologist makes the final call; the system only draws attention.

THE COURSE

What Does the Course Cover?

1. Imaging physics: X-ray, CT, MRI, ultrasound
2. DICOM, NIfTI and medical image data
3. Image reconstruction: filtered back-projection
4. Denoising and image enhancement
5. Classical segmentation: thresholds, regions, level sets
6. Deep segmentation: U-Net and nnU-Net
7. Rigid and deformable registration
8. Atlas-based analysis and volumetry
9. Classification and computer-aided diagnosis
10. Learning from little data: augmentation, transfer
11. Clinical validation: Dice, ROC, reader studies
12. Regulation, privacy and clinical deployment

Skills you will gain

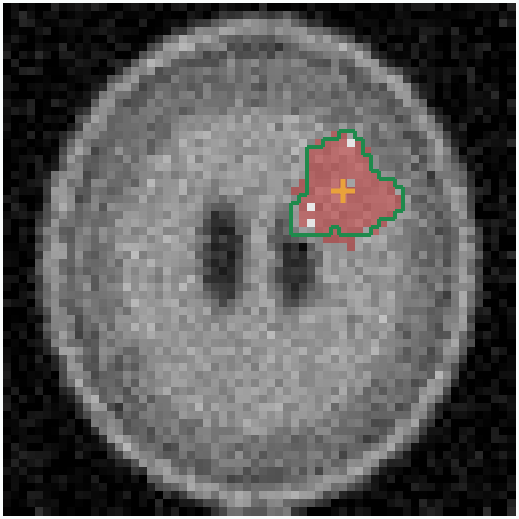
- reading DICOM with pydicom
- annotating images in 3D Slicer
- training segmentation models with MONAI
- registration with SimpleITK
- reporting Dice and ROC results

A typical semester; the order and coverage vary between universities.

TRY IT

Segment the Tumor in an MRI Slice

In a contrast-enhanced brain MRI slice the tumor appears bright; the radiologist wants its volume. Click inside the tumor to place a seed and change the tolerance: does the region cover the tumor or leak into healthy tissue?



Tolerance = ±**30** gray levels

Smooth first | **Expert mask**

Region: 123 pixels = **1,107 mm²**
Expert: 122 pixels = 1,098 mm²
Dice score: **0.89**

■ found region ▭ expert-drawn boundary. Pixel 3×3 mm. The closer **Dice** is to 1, the better the segmentation.

In medical imaging a segmentation becomes a number: volume, growth, response to treatment. The course teaches measuring a method against **expert labels** with Dice, and preprocessing against noise.

CONTEXT

Prerequisites

Courses worth knowing before this one, and why they are needed.

Course	Why it is needed
Image Processing	denoising, segmentation, registration
Computer Vision	lesion detection, classification
Programming	working with multidimensional arrays in Python
Physics	basics of X-ray absorption and magnetic resonance
Intro to anatomy and physiology	recognizing organs and how they appear in images

CONTEXT

Are You Ready for This Course?

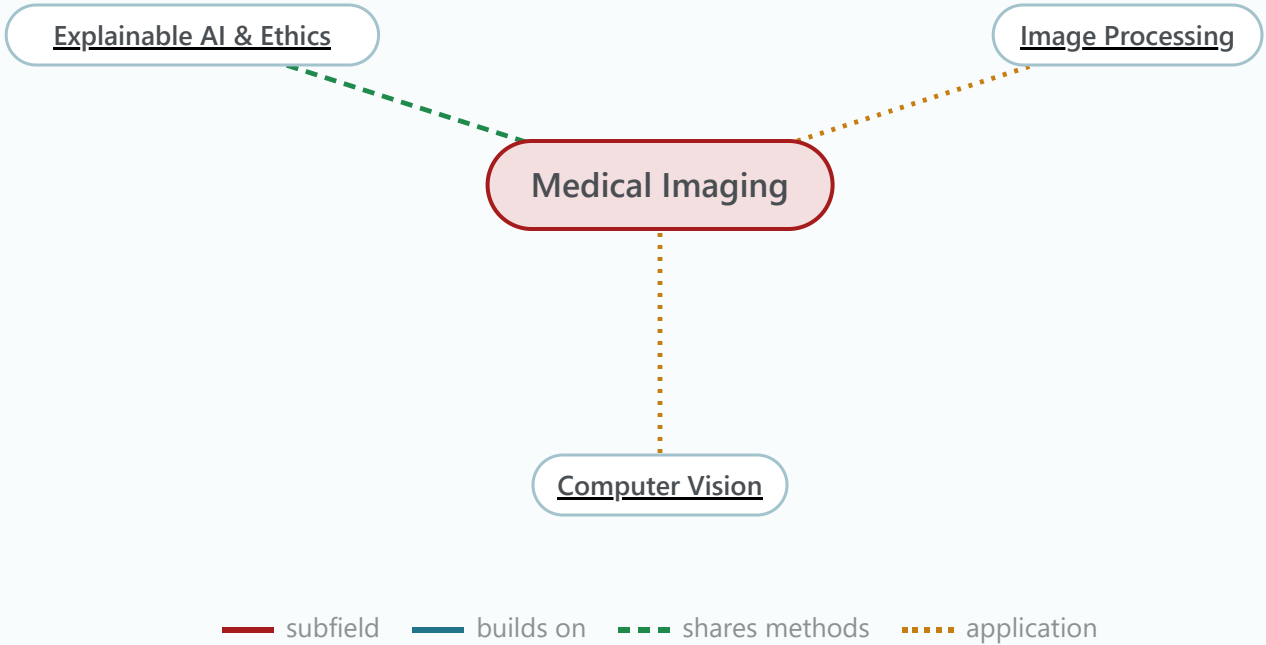
Ask yourself; for every "no", the right-hand column shows where to start.

Ask yourself	If not, start with
Can you apply filtering, thresholding and edge detection to an image?	Image Processing
Can you explain how a convolutional neural network is trained?	Deep Learning
Can you compute and interpret sensitivity and specificity?	Probability and Statistics
Can you slice a three-dimensional NumPy array in Python?	An intro programming course, then a short NumPy tutorial

If you can say "yes" to most of these, you are ready; one or two gaps can be closed during the semester.

CONTEXT

Neighbouring Courses



OVERLAP

Overlapping Courses

Relation	Course	Shared ground
application	This course applies the methods of Image Processing	denoising, segmentation, registration
application	This course applies the methods of Computer Vision	lesion detection, classification
shares methods	Shares methods with Explainable AI and AI Ethics	explainability and accountability in diagnosis

DIFFERENCES

Where It Differs (1/2)

Criterion	Medical Imaging	Image Processing	Computer Vision
Core question	How do X-ray, CT and MRI images help diagnosis and treatment?	How is an image enhanced, transformed and divided into regions?	How does a computer understand what is in an image or a video?
Approach	image processing plus clinical knowledge	pixel-level transformations	extracting meaning from images
Input	CT, MRI, X-ray, ultrasound	image	image or video
Output	segmentation, measurement, diagnostic suggestion	processed image or regions	label, box, depth, 3D structure
Success measure	sensitivity, specificity, Dice	visual quality, PSNR	accuracy, mAP, IoU
Math load	medium-high	medium: convolution, Fourier	high: geometry, linear algebra
Data needs	little; labels are expensive	little; learning not required	many labeled images

- **Image Processing:** Image processing provides the general methods; this course applies them to 3D medical volumes, with device physics and clinical constraints.
- **Computer Vision:** Vision works with millions of labeled photos; in medical imaging labels are scarce, need experts, and errors are costly.

DIFFERENCES

Where It Differs (2/2)

Criterion	Medical Imaging	Explainable AI & Ethics	Bioinformatics
Core question	How do X-ray, CT and MRI images help diagnosis and treatment?	When can we trust an AI system's decision, and who is responsible for the outcome?	How are biological sequences and molecular data analyzed computationally?
Approach	image processing plus clinical knowledge	human-centered evaluation	biology, algorithms and statistics
Input	CT, MRI, X-ray, ultrasound	model, its data and context of use	DNA, protein sequences, expression data
Output	segmentation, measurement, diagnostic suggestion	explanation, audit report, principles	alignment, structure, function prediction
Success measure	sensitivity, specificity, Dice	interpretability, fairness, accountability	biological validity
Math load	medium-high	low-medium	medium-high
Data needs	little; labels are expensive	the model's own data	huge sequence datasets

- **Explainable AI & Ethics:** The ethics course treats explainability in general; this course uses it for physician trust and accountability.
- **Bioinformatics:** Bioinformatics looks at sequence and molecular data; medical image processing looks at images of organs and tissue.

PRACTICE

Applications

Radiology
Finding lung nodules and fractures.

Oncology
Tumor segmentation, radiotherapy planning.

Ophthalmology
Screening retinal images for retinopathy.

Neurology
Lesion and volume measurement in brain MRI.

Pathology
Counting cells in tissue sections.

Surgery
Preoperative 3D models and navigation.

PRACTICE

Risks and Ethics

Shortcut learning

A model can learn scanner labels or hospital markers instead of disease; this has been shown in chest X-ray models.

Failure at another hospital

Performance drops with other scanners, protocols and patients; validation at a single site is not enough.

Re-identification

DICOM headers hold identifiers and faces can be rebuilt from head MRI; anonymization is more than deleting a name.

Automation bias

Clinicians may miss a lesion the model did not flag; a support tool can loosen human scrutiny.

In depth: Explainable AI and AI Ethics.

PRACTICE

Tools and Datasets

Tools and libraries

3D Slicer | ITK / SimpleITK | MONAI | nnU-Net | pydicom | NiBabel | ITK-SNAP

Well-known datasets

BraTS (brain tumor) | LIDC-IDRI (lung CT) | ChestX-ray14 | CheXpert | Medical Segmentation Decathlon

Medical datasets are **personal health data**: they cannot be shared without anonymization, ethics board approval and a data use agreement.

PRACTICE

Common Misconceptions

Common belief	In fact
AI will soon replace radiologists.	Systems support narrow tasks; diagnosis, context and responsibility stay with the physician.
High accuracy means clinical success.	For rare diseases sensitivity and specificity matter; performance may drop at another hospital.
A medical image is an ordinary photo.	CT and MRI are often 3D, and pixel values carry physical meaning (Hounsfield units).
Collecting lots of data is easy.	Labels take expert time, and the data is protected by privacy law.

NEXT STEPS

After This Course

Courses that build on or extend this one

This course is mostly an end point: it builds on the others.

Related careers

Medical imaging engineer

Biomedical software developer

Clinical AI researcher

Medical device regulatory specialist

Medical image processing is the most tightly regulated application of image processing and computer vision. From here you can move on to segmentation with deep learning, to clinical trust with explainable AI, or to the molecular level with bioinformatics.

NEXT STEPS

Further Reading

- Suetens, *Fundamentals of Medical Imaging*, 3rd ed., Cambridge University Press, 2017.
- Bankman (ed.), *Handbook of Medical Image Processing and Analysis*, 2nd ed., Academic Press, 2008.
- Ronneberger, Fischer and Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation", *MICCAI*, 2015.
- Litjens et al., "A Survey on Deep Learning in Medical Image Analysis", *Medical Image Analysis* 42, 2017.

The End

Chapter 20: Medical Image Processing

← Previous chapter · Map of the field · Next chapter →