Image Processing

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

What is an image? | Enhancement | Filters and convolution | Edges and segmentation | Compression | 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_17_Image_Processing.html

What Question Does It Ask?

> **How is an image enhanced, transformed and divided into regions?**

Image processing studies methods that take an image and turn it into **a new image** or into regions: reducing noise, boosting contrast, finding edges and object boundaries. Both input and output are images; understanding what is in the image is the job of computer vision.

At a glance

APPROACH pixel-level transformations	**INPUT** image	**OUTPUT** processed image or regions
SUCCESS MEASURE visual quality, PSNR	**MATH LOAD** medium: convolution, Fourier	**DATA NEEDS** little; learning not required

INTRODUCTION

A Short History

1920s
The Bartlane system sends digitally coded newspaper photos across the Atlantic by cable.

1957
Russell Kirsch makes the first digitally scanned photo: a 176×176-pixel image of his infant son.

1964
JPL corrects Ranger 7's lunar images by computer; digital image processing is born.

1979
Otsu publishes his method for picking a threshold from the histogram; still the most used thresholding.

1986
The Canny edge detector appears: noise-robust, thin and continuous edges.

1992
The JPEG standard is adopted; DCT-based lossy compression becomes the common language of photos.

THE COURSE

Core Concepts

Pixels and channels
An image is a matrix of numbers; color has three channels.

Histogram
Brightness distribution; contrast via histogram equalization.

Spatial filter
Convolution with neighboring pixels: blurring, sharpening.

Frequency domain
2D Fourier transform; filtering periodic noise.

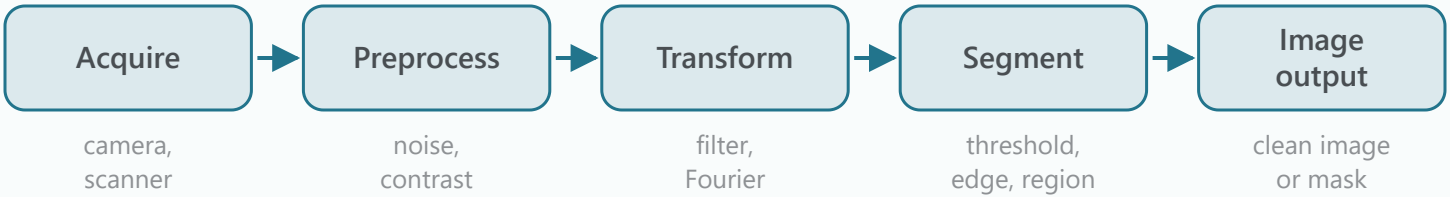
Edge detection
Brightness jumps via gradient, Sobel and Canny.

Thresholding
Separating foreground from background; Otsu's method.

Morphology
Erosion and dilation: cleaning up binary shapes.

Segmentation
Splitting an image into regions: region growing, watershed.

A Typical Workflow



Example: An old scanned document is first cleaned of specks with a median filter, its contrast is boosted by histogram equalization, it is binarized with Otsu's threshold, and the letters are repaired with morphological operations. The output is a clean image; reading the letters is the job of computer vision.

THE COURSE

What Does the Course Cover?

1. Image formation, sampling and quantization
2. Point operations and histogram equalization
3. Spatial filtering: smoothing and sharpening
4. The 2D Fourier transform and frequency-domain filtering
5. Image restoration: noise models, the Wiener filter
6. Colour spaces and colour image processing
7. Geometric transforms and interpolation
8. Edge detection: gradients, Sobel, Canny
9. Thresholding and Otsu's method
10. Morphological operations
11. Segmentation: region growing, watershed, labelling
12. Image compression: the DCT and JPEG

Skills you will gain

building image pipelines with OpenCV | segmentation with scikit-image | designing filters and kernels | interpreting histograms and spectra | measuring objects in ImageJ

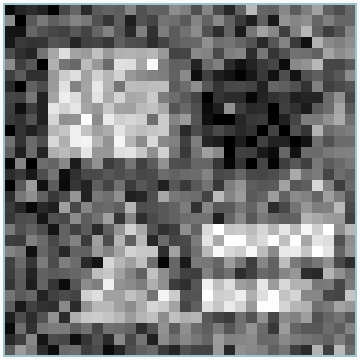
A typical semester; the order and coverage vary between universities.

TRY IT

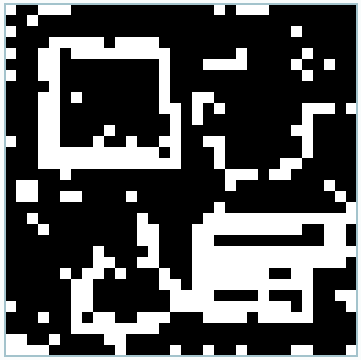
Finding Edges in a Noisy Image

A quality control camera returned a 32×32 grayscale image of a part; the sensor is noisy. Toggle the steps and slide the threshold: how many of the detected edges lie on a real shape boundary?

Input (noisy)



Output: edge strength → threshold



1. Blur | **2. Sobel edges** | **3. Threshold**

Threshold = **70**

Edge pixels: 302 / 1024
On a true boundary: 78%
Boundary captured: 67%
Balance (F1): **0.72**

Each step is a real 3×3 **convolution**. Blur and lower the threshold a little: noise edges drop and F1 rises.

Gauss (÷16)

1	2	1
2	4	2
1	2	1

Sobel Gx

−1	0	1
−2	0	2
−1	0	1

Sobel Gy

−1	−2	−1
0	0	0
1	2	1

Threshold

≥ 70

In image processing both input and output are **images**: filters transform pixels, they do not decide. The course teaches what each kernel does and how the order of steps changes the result.

CONTEXT

Prerequisites

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

Course	Why it is needed
Digital Signal Processing	2D signals, convolution, Fourier transform
Linear algebra and calculus	matrix operations, derivatives and gradients: the language of filters and edges
Programming	loops over 2D arrays in Python or C
Data Structures	queues and union-find: region growing and connected-component labeling

CONTEXT

Are You Ready for This Course?

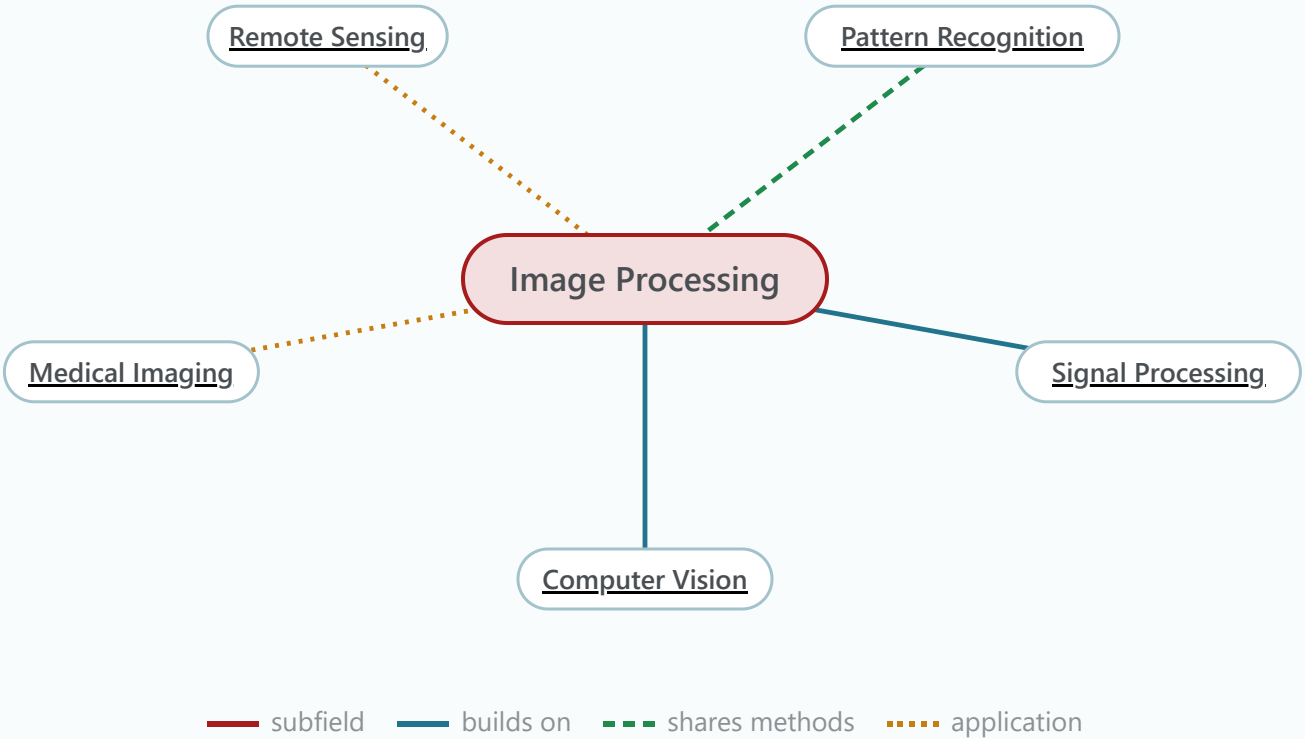
Ask yourself; for every "no", the right-hand column shows where to start.

Ask yourself	If not, start with
Can you visit every element of a matrix with nested loops?	An intro programming course, then a short NumPy tutorial
Can you explain what a partial derivative and a gradient are?	Derivatives and partial derivatives from calculus
Can you multiply matrices and take a dot product?	Linear Algebra
Can you explain what the frequency components of a signal are?	Digital Signal Processing

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
builds on	This course builds on Digital Signal Processing	2D signals, convolution, Fourier transform
builds on	Computer Vision builds on this course	edges, filters, segmentation
application	Medical Image Processing applies the methods of this course	denoising, segmentation, registration
application	Remote Sensing applies the methods of this course	multispectral images, correction
shares methods	Shares methods with Pattern Recognition	feature extraction, template matching

Where It Differs (1/2)

Criterion	Image Processing	Computer Vision	Signal Processing
Core question	How is an image enhanced, transformed and divided into regions?	How does a computer understand what is in an image or a video?	How are sampled signals analyzed, filtered and transformed?
Approach	pixel-level transformations	extracting meaning from images	mathematical transforms and filters
Input	image	image or video	sampled signal
Output	processed image or regions	label, box, depth, 3D structure	filtered signal, spectrum
Success measure	visual quality, PSNR	accuracy, mAP, IoU	signal-to-noise ratio, frequency response
Math load	medium: convolution, Fourier	high: geometry, linear algebra	high: complex numbers, Fourier
Data needs	little; learning not required	many labeled images	no learning; the signal suffices

- **Computer Vision:** In image processing both input and output are images; in computer vision the input is an image and the output is meaning: a label, a box, depth.
- **Signal Processing:** Signal processing works with a 1D time signal; image processing carries the same convolution and Fourier tools to two dimensions.

DIFFERENCES

Where It Differs (2/2)

Criterion	Image Processing	Computer Graphics	Medical Imaging
Core question	How is an image enhanced, transformed and divided into regions?	How is a realistic image produced from the description of a scene?	How do X-ray, CT and MRI images help diagnosis and treatment?
Approach	pixel-level transformations	geometry and light simulation	image processing plus clinical knowledge
Input	image	scene: models, lights, camera	CT, MRI, X-ray, ultrasound
Output	processed image or regions	image or animation	segmentation, measurement, diagnostic suggestion
Success measure	visual quality, PSNR	realism, frame rate	sensitivity, specificity, Dice
Math load	medium: convolution, Fourier	high: linear algebra, geometry	medium-high
Data needs	little; learning not required	none	little; labels are expensive

- **Computer Graphics:** Graphics produces an image from a scene description; image processing takes an existing image and transforms it.
- **Medical Imaging:** Medical image processing applies this course's methods to medical images such as CT and MRI, under clinical constraints.

PRACTICE

Applications

Photo editing
Brightness, sharpness, noise reduction.

Document scanning
Deskewing, converting to black and white.

Compression
JPEG, PNG and video coding.

Industrial inspection
Isolating scratches and stains on products.

Microscopy
Segmenting and counting cells.

Preprocessing
Clean input for vision and medical models.

PRACTICE

Risks and Ethics

Image forgery

The same tools can remove or add objects in a photo, and an edited image can be passed off as evidence.

Corrupted evidence

Lossy JPEG and sharpening change detail in forensic or medical images; the original file must be kept.

Invented detail

Super-resolution and denoising can produce detail that was never in the scene, and it may be taken as real.

Hand-tuned thresholds

Thresholds and parameters picked by looking at one image fail silently when the lighting changes.

In depth: Explainable AI and AI Ethics.

PRACTICE

Tools and Datasets

Tools and libraries

OpenCV | scikit-image | Pillow | ImageJ / Fiji | MATLAB Image Processing Toolbox | GIMP

Classic test images and datasets

USC-SIPI test images | Kodak image set | BSDS500 (Berkeley segmentation) | DIV2K (super-resolution) | Cameraman and Peppers

In image processing **a single image** is often enough: the method is defined in advance, not learned from data. Large labeled datasets are what computer vision needs.

PRACTICE

Common Misconceptions

Common belief	In fact
Image processing and computer vision are the same thing.	Image processing transforms the image; vision extracts meaning from it. One is a step before the other.
Sharpening brings back lost detail.	Sharpening accentuates existing edges; detail that was never recorded does not come back.
Deep learning does everything now.	Classical methods are still essential in calibration, compression, preprocessing and fast industrial inspection.
Enlarging an image increases its resolution.	Interpolation creates new pixels but no new information; the image just gets bigger.

NEXT STEPS

After This Course

Courses that build on or extend this one

Computer Vision | Medical Image Processing | Remote Sensing

Related careers

Image processing engineer | Embedded vision engineer | Quality control systems engineer | Biomedical imaging specialist

Image processing is the common first step for every course that works with images. To extract meaning from images, move on to computer vision; for applications, to medical image processing or remote sensing. In the opposite direction, computer graphics studies how to produce images.

NEXT STEPS

Further Reading

- Gonzalez and Woods, *Digital Image Processing*, 4th ed., Pearson, 2018.
- Jain, *Fundamentals of Digital Image Processing*, Prentice Hall, 1989.
- Burger and Burge, *Digital Image Processing: An Algorithmic Introduction*, 3rd ed., Springer, 2022.
- Canny, "A Computational Approach to Edge Detection", *IEEE Transactions on Pattern Analysis and Machine Intelligence* 8(6), 1986.

The End

Chapter 17: Image Processing

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