WORKING WITH AI · IMAGE ANALYSIS · SESSION 3

Segmentation & Colocalization

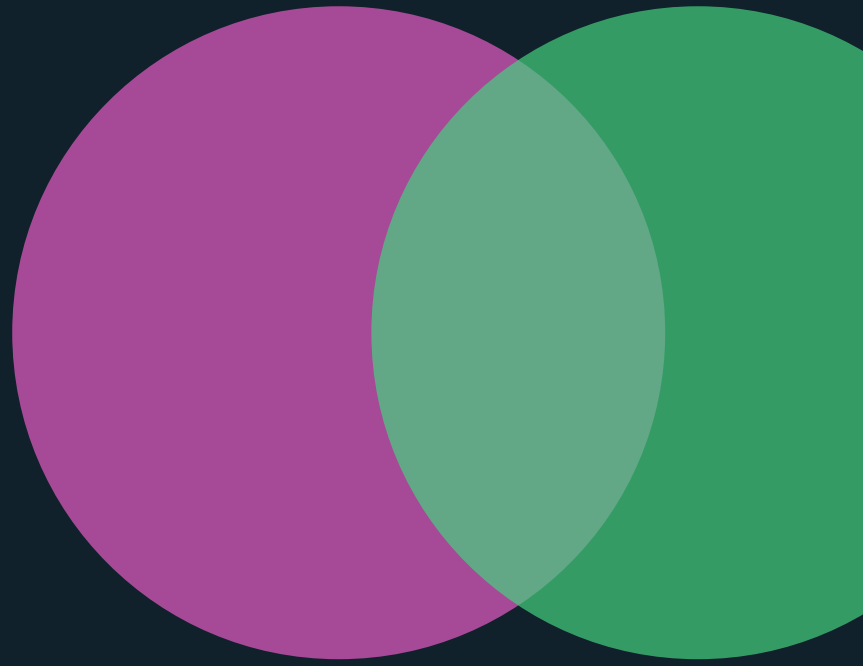

Two signals — what belongs together, and how do we prove it?

Build a colocalization analyzer, live · then turn it loose on your own data

TODAY · ~55 MINUTES

How the hour runs

1

15 min

Your builds

Three of you present the tools you made — 5 minutes each

2

15 min

Foundations

The statistics of two signals: segmentation, colocalization, spacing

3

20 min

Build it live

We construct the analyzer together, one rung at a time

4

5 min

Discovery

Detection becomes a new question — on your own images

Reproducibility · Accountability · Transparency · Verification — the four habits run through every part.

PART 1 · 15 MINUTES

Your builds

Three presenters · five minutes each · then we move on

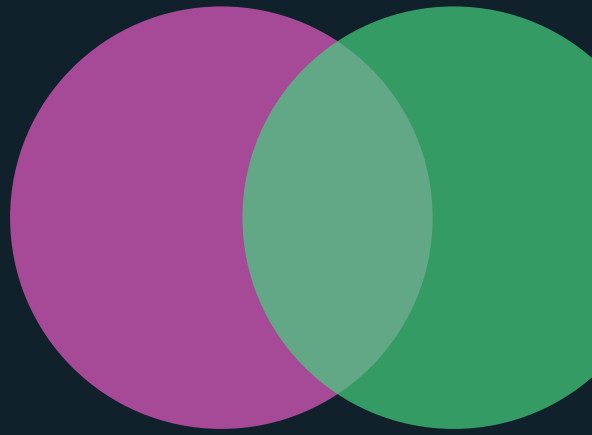

PART 1 · SHOWCASE

As you watch, ask two questions

What did they measure?

Which property of the image did they turn into a number — and why that one?

How did they verify it?

What convinced them the tool was right? One image is never enough.

Presenters: keep it to 5 minutes — the description you gave Claude, the tool, and one thing that surprised you.

PART 2 · 10–15 MINUTES

The two-signal question

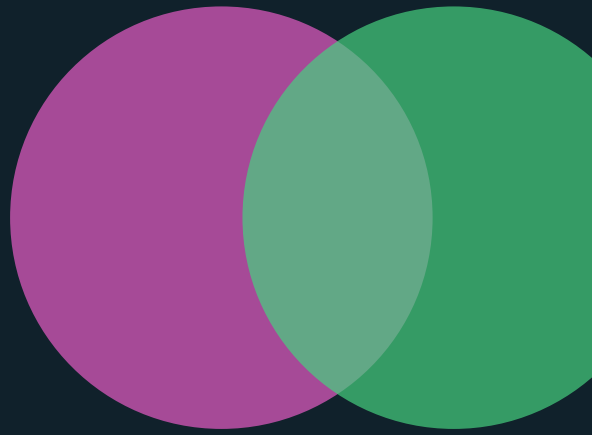

You have two labels. Do they occupy the same places — and how would you ever know?

From an image to an answer, in three moves

1

Segment each channel

Find the individual objects (puncta) in each image separately.

2

Describe the objects

Count, size, and spacing — per channel, in real units.

3

Relate the two

Do the objects of one channel coincide with the other — and is that more than chance?

Today's tool does all three — and refuses to skip the third.

Turning brightness into objects

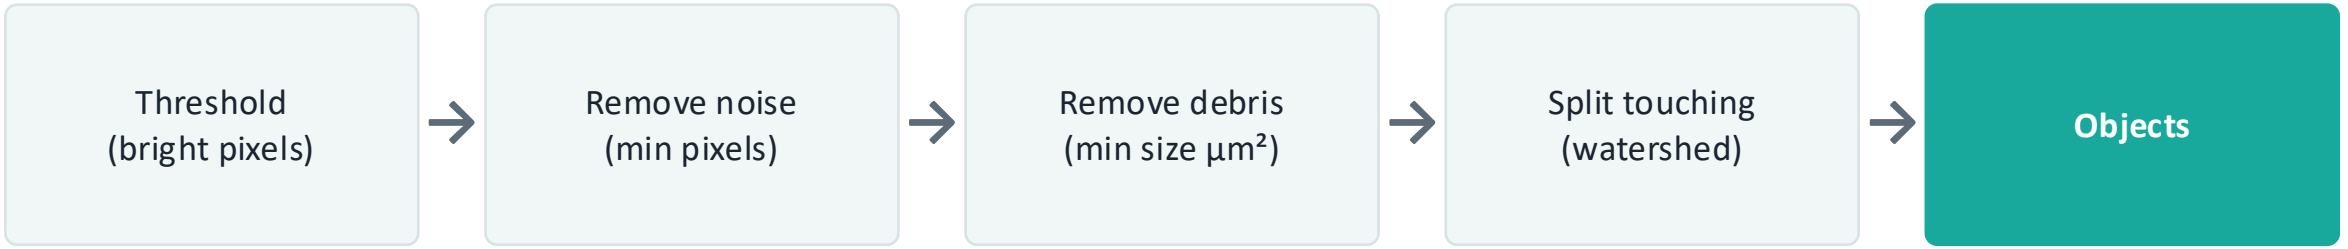

The overlay is the interface between the code and the biology.

You don't need to read the code. If the outline lands on what a biologist would call an object, the segmentation is right — regardless of how it's written. If it doesn't, no amount of elegant code saves it. Each control is a real decision you own: how bright, how big, how separate.

Colocalization is not one number

Pearson

Do the two intensities rise and fall together? −1 to +1.

Manders M1 / M2

What fraction of each channel overlaps the other? Directional.

Object-based

For each object, is a partner within X µm? A % with partners.

Significance

Is the overlap more than chance? A randomization test.

Coloc map

Where does overlap happen? A pixel-by-pixel picture.

Each answers a different question. Report several — they disagree in informative ways.

FOUNDATIONS · PEARSON

The most common number — and its trap

Pearson asks whether the two channels brighten together. Simple, standard — and easy to inflate.

WHOLE IMAGE

0.72

"Strongly colocalized!" But most of the image is shared black background — co-absence masquerading as correlation.

SIGNAL ONLY (THRESHOLDED)

0.18

Restricted to where signal actually is, the honest answer: the two channels barely co-vary. Same image, opposite conclusion.

Lesson: measure colocalization on the segmented signal, never the whole frame.

Overlap is directional

M1

Of all the Channel 1 signal, what fraction sits on Channel 2?

M2

Of all the Channel 2 signal, what fraction sits on Channel 1?

Why they differ

A few large Channel-1 blobs can each sit on a Channel-2 punctum (high M1) while covering only a sliver of the abundant Channel-2 signal (low M2).

Manders is not one "% overlap." The two numbers together describe an asymmetric relationship — which marker is the subset of which.

From pixels to objects — and where they sit

Object-based colocalization

Object has a partner within X µm → colocalized

No partner nearby → not colocalized

Counts objects, not pixels — robust to brightness, close to how you'd score it by eye.

Nearest-neighbour & arrangement

How far is each object from its closest same-type neighbour? The mean distance, plus a clustered–random–regular score (Clark-Evans):

< 1	≈ 1	> 1
clustered	random	evenly spaced

Two kinds of randomness: sizes can be random while positions are orderly.

Overlap alone proves nothing

Two busy channels will overlap somewhat by pure chance. The only honest test: compare what you see to what randomness gives.

1

Observed

Measure the real overlap in your image.

2

Shuffle

Randomly reposition the objects, hundreds of times.

3

Chance distribution

How much overlap does randomness usually produce?

4

p-value

How surprising is the real value against that backdrop?

A weak-but-significant result and a strong-but-chance result are both real findings. The test tells them apart.

A picture of colocalization

Li's Intensity Correlation Analysis

For every pixel, multiply each channel's distance from its own average. Sum those products and you have Pearson — so the map is Pearson, taken apart pixel by pixel.

mutually exclusive — colocalized

Why a map matters

A single coefficient hides structure. The map shows whether overlap is everywhere, or concentrated in a few hotspots while the rest of the cell keeps the channels apart.

Warm where the two rise together. Cool where one is bright and the other dark. Computed on signal only — background can't vote.

FOUNDATIONS · HONESTY

What the numbers can't tell you

1	**Resolution**	Two dyes within ~250 nm look coincident even when they're on separate molecules. Overlap ≠ contact.
2	**Diffraction limit**	Object sizes are the microscope's blur, not the true size, until you go to super-resolution.
3	**Coincidence ≠ interaction**	Colocalization is spatial overlap. It suggests, it never proves, a molecular relationship.

A good tool reports these limits itself. Ours does — in plain English.

PART 3 · 20 MINUTES

Build it — together

Screen-share. Nine prompts, one rung at a time. You describe the science; Claude writes the code; we watch each step land.

Nine sends, R0 → R7

R0	Describe the two channels — no building yet
R1	Load & display · the naive whole-image Pearson (the trap)
R2	Segment: threshold → noise floor → size → split touching
R3	Measure objects: count, size, spacing per channel
R4	Colocalize three ways: signal Pearson · Manders · object overlay
R5	Ask Claude how to test it — then implement the randomization
R6	The colocalization map (Li's ICA)
R7	Quality control, scale bar & a plain-English summary
R8	Validate on known ground truth — precision/recall & a report

Highlighted rungs are the turns — where a number changes, or a decision becomes yours.

PART 3 · WHILE IT BUILDS

Three moments to watch for

The Pearson drop

When R4 restricts to signal, the number falls — 0.72 → 0.18. Same image, honest answer.

You ask before you test

At R5 you ask Claude how to prove the result, and reason about the method before any code is written.

A negative result stands

If the two channels separate, 'not above chance' is the finding. The tool isn't broken — the biology is the answer.

PART 4 · 5 MINUTES

Detection becomes discovery

The tool answered 'do they overlap?' The better question is the one it makes you ask next.

PART 4 · DISCOVERY

Where this goes on your own images

Run it on your data

Load your two channels. Let the summary tell you what the image can — and can't — support.

Ask about pattern

Beyond overlap: are objects clustered? Ripley's K and pair-correlation quantify spatial organization.

Go below the limit

Colocalization at the diffraction limit is a hypothesis. Super-resolution tests it — that's Session 4, with Manuel Navedo (STORM).

The ceiling on what you can measure is now a scientific choice, not a software limit.

TAKE IT HOME

Four tools, one habit

- **Segment before you relate** — objects first, relationships second
- **Measure on signal, not background** — or Pearson lies to you
- **Test against chance** — overlap alone proves nothing
- **Report what you can't see** — resolution is part of the result

NEXT

Session 4

The Analytical Frontier — super-resolution (STORM) with Manuel Navedo. Below the diffraction limit.

Homework: build the analyzer yourself from the prompt sheet — then run it on one of your own image pairs and bring the summary.