WORKING WITH AI · SESSION 2 HOMEWORK

Build Your Own Analysis Tool

A prompt-by-prompt guide — then discovery on your own

Today you ran tools other people built. This homework is about building one yourself, with Claude, and then using that skill to chase a question nobody handed you. **Part A is fully scaffolded** — copy the prompts, get a working tool, confirm it against an answer you can check. **Part B is yours** — same method, a metric you choose, a hypothesis you didn't have before.

WHAT YOU'LL NEED — three files from today

Claude (the same access you used in the session), plus the practice linescan package: Session2_Practice_Linescan.tif (the image you'll analyze), Session2_Practice_Linescan_metadata.txt (the acquisition numbers you'll put into your description), and Session2_Practice_Linescan_GroundTruth.csv (the answer key — keep it closed until you check your result). No coding required: you describe, Claude writes the code, you verify.

HOW TO USE THE THREE FILES

Upload the .tif to Claude when you're ready to analyze it. **Open the metadata .txt yourself and read its numbers into your description prompt** (500 lines/s, 0.30 µm/px, 256×750, raw F, F_0 = first 100 ms) — that's the "rich description with metadata" fundamental from session one, and it's why the tool won't have to guess. **Do NOT open or upload the ground-truth key until the verify step**

WANT TO REBUILD THE WHOLE DETECTOR?

Everything you watched — all nine rungs — is in **Session2_Student_Prompts_Rebuild_the_Detector.docx**. The exact prompts, in order, nothing hidden. It took 24 minutes in the room. Claude plants its own events, so your numbers will differ from ours — that's correct, and it's the point: you own the ground truth, so you can score the detector against it. Prompt 10 is the capstone: it teaches your tool to load the practice TIFF, and asks Claude what the tool can no longer verify once it's looking at a file it didn't plant. Parts A and B below are the shorter path if you'd rather practice the loop first.

PART A Build a tool for a known metric

The metric is prescribed and the answer is known, on purpose — so you can focus on learning the loop, not on picking a metric or wondering whether you got it right. The loop is three moves: **describe, build, verify.**

Prompt 1 — Describe (cover all four dimensions)

This is the prompt that everything else depends on. Fill every bracket. If you leave one vague, Claude will guess — and the four dimensions from session one are exactly the four things it shouldn't have to guess. Don't skip any, even if it feels obvious.

COPY, FILL THE BRACKETS, SEND TO CLAUDE:

```
I'm going to have you build a small analysis tool. First, here is the full
context for the data, in four dimensions:

SAMPLE: [organism, cell type, preparation, condition — e.g. adult mouse
  ventricular myocyte, resting, Fluo-4-loaded]

IMAGING: [modality, pixel size, bit depth, channel, sampling rate — e.g.
  confocal linescan, 500 lines/s, 0.30 um/pixel, raw fluorescence F, single
  Fluo-4 channel]

STRUCTURE: [what the events look like in space and time, and how to tell
  them apart — e.g. brief local delta-F/F0 rises are sparks; one slower,
  whole-cell-wide rise is a transient, which is NOT a spark]

QUESTION: [the biological question — e.g. is the timing of sparks regular
  or irregular? What are the intervals between consecutive sparks?]
```

```
Please read this back to me in your own words and tell me if anything is
ambiguous or missing before we go further.
```

Why the last line matters: asking Claude to reflect the description back is a free check on your own clarity. If it misreads something, your description had a gap — fix it here, before any code exists.

Prompt 2 — Build the tool

Now ask for the tool itself. Notice this prompt tells Claude to show its work — detected spark times first — so you can catch an error before it's buried inside a summary number.

SEND THIS NEXT:

```
Build me a self-contained HTML page that measures INTER-SPARK INTERVALS
for a linescan like the one I described. It should:

1. Let me enter the time of each detected spark (in milliseconds).
2. Compute the interval between each pair of consecutive sparks.
3. Report every interval, plus the mean, standard deviation, and
   coefficient of variation (SD / mean).
4. Show the spark times I entered back to me first, so I can confirm
   them before trusting the intervals.
5. Do NOT treat the two whole-cell transients as sparks.

Keep it simple and readable. No login, no external libraries.
```

You'll get a working page you can open in a browser. If something looks off, just tell Claude what you see — "the intervals are blank," "it counted four events" — and let it revise. That back-and-forth is normal, and it's the skill.

Prompt 3 — Get the spark times from the image, then verify

Your Part A tool takes spark times as input. Two honest ways to get them from the actual practice image — pick either:

Option 1 (recommended) — analyze the TIFF. Upload Session2_Practice_Linescan.tif to Claude and ask it to detect the spark times for you, stating the calibration from the metadata: "This is a 500 lines/s linescan (2 ms/column). Find the sparks and give me their peak times in milliseconds — but not the two whole-cell transients." Then feed those times into your interval tool.

Option 2 — read them off in ImageJ/FIJI. Open the .tif, hover over each spark to read its column, and multiply by 2 ms. This is the manual path, and it makes the column-to-time conversion concrete.

Now open the answer key. Only at this point, open Session2_Practice_Linescan_GroundTruth.csv. It lists the five planted sparks (peak times 490, 650, 850, 1030, 1170 ms), the two transients, and the four intervals your tool should produce. Compare — did your analysis of the image recover them?

THE ANSWER YOUR TOOL MUST PRODUCE

Four intervals: 160, 200, 180 and 140 ms. If your tool says anything else, the tool is wrong — the linescan is fixed and known. The most common bug: entering column numbers instead of milliseconds (remember, 1 column = 2 ms at 500 lines/s), which gives you exactly half the right answer.

If it doesn't match, hand the discrepancy back to Claude:

IF THE NUMBERS ARE WRONG:

```
My tool returned [what you got], but the known intervals for these spark
times should be 160, 200, 180 and 140 ms. Here are the spark times I
entered: 490, 650, 850, 1030, 1170 ms. Find and fix the error, and
explain what went wrong.
```

THE PART EVERYONE MISSES — AND THAT'S THE POINT

Your tool also reports a coefficient of variation — it should land near 0.15. Four intervals beats two, but it is still a very small n: a CV from four gaps carries enormous uncertainty, and you cannot really characterize the regularity of a process from it.

Report it in a paper without saying so and a reviewer will (rightly) object. Recognizing when NOT to trust a number your own tool produced is the real deliverable of Part A.

PART B Now do it on your own — for discovery

Part A taught you the loop on a question with a known answer. Part B removes the training wheels. You've now built one tool; build a second one for a metric you choose — one that could tell you something the conventional measures (amplitude, frequency, kinetics) can't.

Pick one of the discovery-menu approaches from today's session — or invent your own:

- **Coherence —** are spark positions along the cell clustered, random, or evenly spaced?
- **Entropy —** how complex or ordered is the $\Delta F/F_0$ trace across timescales?
- **Site reutilization —** do the same locations fire repeatedly, or does activity spread across the cell?
- **Transient uniformity —** does the whole cell rise together, or does a wavefront propagate?

Then follow the exact same three moves you just practiced — this time writing the prompts yourself:

- **DESCRIBE —** the same four dimensions. This never changes. If you use your own lab data, describe it just as carefully.
- **BUILD —** ask Claude for a tool that measures your chosen metric, and ask it to show intermediate results you can sanity-check.
- **VERIFY —** this is the hard, important part when there's no answer key. Ask Claude to generate synthetic data with a known pattern (e.g. “make a linescan where sparks are deliberately clustered”), run your tool on it, and confirm it detects what you planted. A tool you haven't tested on known input is a tool you can't trust.

WHAT TO TURN IN

Three things: (1) your working tool from Part A, verified at 160 / 200 / 180 / 140 ms; (2) your own tool from Part B and the metric it measures; and (3) — most important — one or two sentences stating a **hypothesis** your Part B metric let you form. Not a measurement. A question or claim about the biology that you didn't have before you built the tool. That leap — from a detected number to a new idea — is the whole point of this session, and of using AI for discovery at all.

That's the same move that took the field from one observed calcium spark to the local control theory of the heart — a detection became a question, and the question became a theory. You just did the first step of it in one evening.

WHILE YOU'RE AT IT — TWO APPS TO EXPLORE

You saw these demoed in class; now they're yours to break. **Nyquist Explorer** — drag the rate and pixel-size sliders and watch a true signal degrade as sampling coarsens; find the point where your own recordings would start losing the peak. **Fit-Quality Explorer** — flip between the single-exponential, bi-exponential, and model-free tabs on a decay that genuinely has two components, and watch a healthy-looking R^2 fail the residual and baseline checks. Five minutes each; both reward poking at the edge cases.