WORKING WITH AI · REFERENCE CARD

Session 2 — Calcium Signals

Events in Time · builds directly on Session 1 — keep both

THE ONE IDEA

A number you look up isn't automatically the number that applies to your measurement. An indicator's Kd, a sampling rate, a threshold — each has a context it was measured or chosen under. Check that context before you trust the number.

How today worked

talk to the model → build the analysis → validate it. Same arc every session. Today you watched it happen live: we built a calcium event detector with Claude rung by rung, the code getting smarter and the tool growing new panels at each step. You rebuild it yourself, unhurried — the prompts are all in your homework pack.

Four tools ship with today's materials — the Nyquist and Fit-Quality explorers, the threshold widget, and the full eight-rung detector we built. All run in any browser, offline, no login. See "Take these home" at the end of this card for what each one is for — and why it's worth the ten minutes.

Imaging, extended for time

The four dimensions from session one still apply in full. Imaging gains a temporal axis — four additions, every time you describe a recording.

Addition	What it covers	Thin → rich
Sampling rate	Lines/s or fps — sets your temporal Nyquist.	"fast enough" → "500 lines/s, 2 ms/line"
Signal type	Raw F, or already computed $\Delta F/F_0$?	"the trace" → "raw F; $\Delta F/F_0$ computed from a stated baseline"
Cell state	Resting, paced, or field-stimulated.	"a cell" → "resting, saponin-permeabilized"
F_0 window	The baseline segment $\Delta F/F_0$ is measured against.	"the start" → "first 100 ms, before any event begins"

The indicator vocabulary (six things to name)

Brightness — = ε (extinction coefficient) × Q (quantum efficiency). Fluo-4's ε ≈ 100,000 $cm^{-1}M^{-1}$ — well above Fura-2's ≈ 34,000.

Apparent Kd — where the F-vs-Ca curve is steepest. Fluo-4 in vitro ≈ 345 nM; in situ ≈ 1,000–1,200 nM — roughly 3× higher, cuvette vs. cell.

Toxicity / buffering — enough dye to see it, not so much it buffers the signal away. SM2's own protocol: ≈10 µM Fluo-4-AM for intact cells.

Photobleaching — image the same field twice before an event starts. A drop with no event is bleaching, not biology.

Dynamic range — Fmax/Fmin — Fluo-4 exceeds 100-fold. A floor-to-ceiling number, not a promise at any one $[Ca^{2+}]$.

Ratiometric vs. single-wavelength — Fura-2 — self-calibrating, slower. Fluo-4 (today's dye) — brighter, faster, no built-in denominator.

No indicator is linear — this is universal, not a Fluo-4 quirk

$F = F_{min} + (F_{max} - F_{min}) \times [Ca^{2+}] / ([Ca^{2+}] + K_d)$ — Grynkiewicz, Poenie & Tsien, 1985. Flat at low calcium, flat at saturation, steepest at Kd. The same $\Delta F/F_0$ means a different $\Delta[Ca^{2+}]$ depending on where you sit on this curve.

The $\Delta F/F_0$ shorthand itself traces to Cheng, Lederer & Cannell, 1993 — the paper that named the calcium spark.

Setting your sampling — worked in today's data

A FLUO-4 SPARK

Rise ~10 ms. Want ≥ 5 points.

interval ≤ 10 ÷ 5 = 2 ms

rate ≥ 500 lines/s

Today's linescan runs at exactly 500 — the floor, not a cushion.

TODAY'S WHOLE-CELL TRANSIENT

Rise ~50 ms. Even 5 points is enough.

interval ≤ 50 ÷ 5 = 10 ms

rate ≥ 100 lines/s

At 500 lines/s that's 25 points — comfortably oversampled.

Same recording, two different margins — know which one you're trusting before you report a time-to-peak. Same check applies spatially: FWHM ÷ pixel size ≥ ~2.3 points (Nyquist).

Trusting a result, extended

The overlay is still the interface. — now it's a true contour, not a box. An elongated or dumbbell-shaped outline is real information — usually two nearby events merged into one.

The $\Delta F/F_0$ event cut is the same threshold decision, in time. — cutoff = background mean + k×SD, set from the background, never auto-adjusted per recording.

Not every real event clears your cutoff. — in a realistic two-population simulation, about 1 in 4 true events fall below a mean+3SD cutoff. That's the shape of the true distribution, not a broken threshold. And it isn't only a quiet-recording problem — a noisier recording hides the same small events, because elevated biological noise raises the floor they have to clear.

A fit produces a number even when it's wrong. — R^2 can sit at 0.97 while a single-exponential fit is systematically wrong. Check three things, not just R^2: residual pattern (a curve or same-sign runs = wrong model), fitted baseline (should sit near zero), and the tail of the overlay. When unsure, fall back to model-free measures — time-to-peak, time-to-50%-decay — less information, never confidently wrong.

Known ground truth, same proof. — the practice linescan has exactly 3 calcium sparks + 1 whole-cell transient planted in it, at known positions and amplitudes.

The detector we built — the nine rungs, to rebuild at home

The exact prompts are in **Session2_Student_Prompts_Rebuild_the_Detector.docx** — the same ones you watched, in order, nothing hidden. You'll generate your own recording, so your numbers will differ from the room's; what you own is the ground truth.

Each rung is one prompt, one owned decision, one fundamental. This is the skeleton of the homework — follow it in your own conversation with Claude.

0 · Describe. — the four dimensions, every number stated. No analysis instructions yet — just context.

1 · Open & display. — generate the linescan, choose your F_0 window (before any event), compute $\Delta F/F_0$, render the heatmap. F_0 is a decision you own. You have this page: Session2_R1_Starting_Point.html — and the one prompt that produced it is archived in a comment at the top of the file.

2 · Naive detection. — one global threshold, no other criteria. It over-counts wildly — a few hundred “events” against 5 real ones. This is supposed to fail.

3 · Real detection. — three separate gates you own — amplitude, duration, width. Each removes a different kind of false positive: the flood → ~14 (faint noise gone) → ~9 (brief flickers gone) → ~6 (narrow streaks gone). Your exact numbers depend on the recording Claude generates; the shape is what matters.

4 · Nyquist check. — per event, measured from its own trace — the 10–90% rise becomes a bandwidth ($f_{bw} = 0.35/\text{rise}$), compared to Nyquist as a continuous margin (no hard cutoff). Everything comes back resolved, and that is the right answer: 500 lines/s genuinely captures a 10 ms rise. Sparks are the tightest (~5–6×), the transient ~13–26×. Anchored to SparkMaster 2 (Tomek et al. 2023, spark TTP ~15 ms → ~7.5×). The point is not alarm — it is that your rate is now justified by a number.

5 · Test. — ask Claude how to test and verify — don't hand it a recipe. It proposes known ground truth + precision/recall. Then implement: reveal the planted events and score. Expect every real event found AND a couple of false positives — signal shoulders that pass all three gates. That honest score beats a perfect one.

6 · Metrics. — amplitude, time-to-peak, FDHM, FWHM, decay tau, signal mass, spark & transient frequency. Amplitude matches what's planted; TTP ~8–10 ms against a planted 10; FWHM straddles the theoretical 4.24 µm. Scatter, no bias — exactly what rung 4 predicted.

7 · Waves. — a propagating event drifts across the cell — a diagonal streak. Detect the drift, measure its slope as a velocity. A new measurable that becomes a hypothesis.

Detection is the floor — discovery is the point

A detection tool answers questions you already knew to ask. The leap that matters is from a detected event to a new hypothesis. **One spark, ten years:** Niggli & Lederer proposed local control (1990), Stern modeled it (1992), Cheng/Lederer/Cannell imaged sparks (1993), Cannell/Cheng/Lederer confirmed the mechanism (1995), Stern's group refined it (1999) — a decade from one observation to a theory that changed how the field sees the heart. SM2 does the first step. AI can help with the other nine years.

Beyond amplitude and frequency — six questions, each ending in a hypothesis: stochasticity (are intervals memoryless?) · entropy (how ordered is the trace?) · spark rate as a precursor · spatial coherence (clustered or random?) · site reutilization (which sites are

"hot"?) · transient uniformity (does the cell rise together?). Open a discovery dialogue: rich SHOW, open WONDER, then PROBE what surprises you.

Before any result goes in a figure

☐ Sampling rate justified against each event's measured rise (Nyquist margin, spatial and temporal) — compared to published kinetics, not an arbitrary cutoff

☐ Indicator's apparent Kd checked against expected resting/peak [Ca^{2+}] — and which condition (in vitro or in situ) the number applies to

☐ $\Delta F/F_0$ threshold set from the background, not eyeballed — never auto-adjusted per recording in a batch

☐ Contour overlay examined and biologically sensible; merged/irregular shapes investigated, not ignored

☐ Any curve fit checked by residuals + baseline, not R^2 alone; model-free fallback used when unsure

☐ Validated against known ground truth

☐ Batch consistency clause included verbatim in the prompt; a quality-control flag is present

Take these home — five tools, five ideas

All four run in any browser: double-click, no login, no install, works offline. Each one exists to make a single idea impossible to un-see. The first two are the ones to open tonight.

Nyquist Explorer — "See it, don't just compute it." One true, continuous spark sits underneath; drag the sampling rate and watch what you would actually have measured at that setting. THE IDEA: sampling rate is not a microscope setting — it decides whether your kinetics are real. WHY IT MATTERS: undersample and you will still get a time-to-peak. It will just be a number about your scan rate, reported with full confidence as a number about the cell.

Fit-Quality Explorer — "Fitting produces a number even when it's wrong." A transient decay with genuinely two kinetic components underneath — fast SR reuptake, slower removal — the way real whole-cell decays usually are. Fit it with a single exponential: R^2 comes back around 0.95–0.97 and the fit is still systematically wrong. THE IDEA: goodness-of-fit does not tell you the model is right. WHY IT MATTERS: R^2 is the number most often quoted and least often earned. Check the residuals for structure, check the fitted baseline sits near zero, check the tail. Each new decay is randomly generated — run it a few times and watch a "good" R^2 keep hiding the same wrong model.

Threshold Widget — "Every cut has a cost." Planted sparks, known answer. Move the threshold: the recording underneath never changes — only what counts as an event does. THE IDEA: a threshold is a trade-off you own, not a right answer you find. WHY IT MATTERS: too high and you lose real events; too low and you count noise. You will be asked to defend where you stood.

R1 Starting Point — The page today's build started from — the synthetic recording, the F_0 selector, the heatmap, and nothing else. THE IDEA: the one prompt that produced it is archived in a comment at the top of the file. WHY IT MATTERS: a tool without its prompt is not reproducible. Open the file, read the prompt, run it yourself — that is the whole loop in one artifact.

Reference Detector — The full eight-rung tool from today's build, with clickable rungs so you can step through the whole climb — bare heatmap to naive flood to gated detection to Nyquist, verification, metrics, waves. THE IDEA: watch the code and the interface mature together. WHY IT MATTERS: it is the answer key for your homework — so build yours first, then open this one to compare.

TODAY'S PORTFOLIO PIECE	WHAT SESSION 3 ASSUMES YOU HAVE
The batch-consistent execution prompt you drafted — DESCRIBE / SPECIFY / VERIFY, with the consistency clause verbatim. Plus tonight's build-your-own homework: a tool you build with Claude, then a hypothesis you form from it. Second piece of your four-session portfolio; the first was the Phone Call Test from session one.	Describe in time · name an indicator's Kd · threshold as a decision · read a contour · validate against ground truth · identical batch parameters, flagged not dropped · open a discovery dialogue and turn a detected pattern into a hypothesis.

Reference: SparkMaster 2 spark kinetics — Tomek J, Nieves-Cintron M, Navedo MF, Ko CY, Bers DM. SparkMaster 2: a new software for automatic analysis of calcium spark data. Circulation Research 2023;133(6):450–462. doi:10.1161/CIRCRESAHA.123.322847