ANTHROP\C

Automated processing of raw NMR and LC-MS data with Claude Opus 5 on Claude Science

Published

Aug 18, 2026

Authors

David Kamber[1]
Claude Science

[1] dkamber@anthropic.com

Summary

Chemists learn whether they made the compound they intended to, and how pure it is, from NMR and LC-MS. Both measurements begin as raw proprietary instrument files: an NMR free induction decay (FID), and an LC-MS binary run file. Turning such files into a phased spectrum, a peak table or a purity figure has been manual work, done in that software, one dataset at a time. That processing recurs with every sample, stands between each experiment and subsequent iteration, and carries a cost that grows with the number of samples.

This write-up follows one routine QC sample from a contract research organization (CRO) through Claude Science, running Claude Opus 5. Claude Science was given the raw files and prompts of one to three short sentences. The ^{1}H NMR data (Part 1) and the LC-MS run (Part 2) were handled in two independent sessions. The two sessions ran side by side, and because each analysis needs only its own file and prompt, the same procedure should extend to many NMR and LC-MS datasets processed concurrently. Each session returned a processed, interpreted result alongside the validation checks supporting it. In each part, Claude Science's output is judged against the CRO's own processing of the same acquisition, preserved in its audit trails and analysis report.

^{1}H NMR: from raw FID to peak table, and the D_2O experiment Claude Science asked for

The CRO acquired the ^{1}H spectrum (32 scans, DMSO-d_6, 400 MHz) on July 31. Its audit trail shows an operator opening the dataset in the vendor's processing software, 11 min after the acquisition ended, and spending 1 min 49 s on phasing, baseline correction, peak picking and twenty integral regions; a D_2O/d-TFA exchange spectrum followed three days later on a second instrument and was processed the same way (16 min after acquisition, 2 min 1 s of commands). Claude Science received the first FID with its acqus header and Prompt 1. Twenty-three minutes later it returned a processed spectrum (Figure 1), an 18-signal peak table and the recommendation to run a D_2O shake; given the exchange FID (Prompt 2), it identified the solvent system, quantified what had exchanged and corrected its own first reading after an internal audit (Figure 2). Where its peak table and the operator's integral regions can be matched, they agree within 0.08 H on a common scale (four weak operator regions of 0.1 to 0.2 H have no counterpart in the table), and its account of the exchange matches the operator's region integrals.

PROMPT 1 · ATTACHMENTS: FID, ACQUS

i have a raw 1H FID: process it: FT, phase, baseline-correct. show me the spectrum. then pick peaks and integrate: give me a table with δ (ppm), multiplicity, J (Hz), and integral.

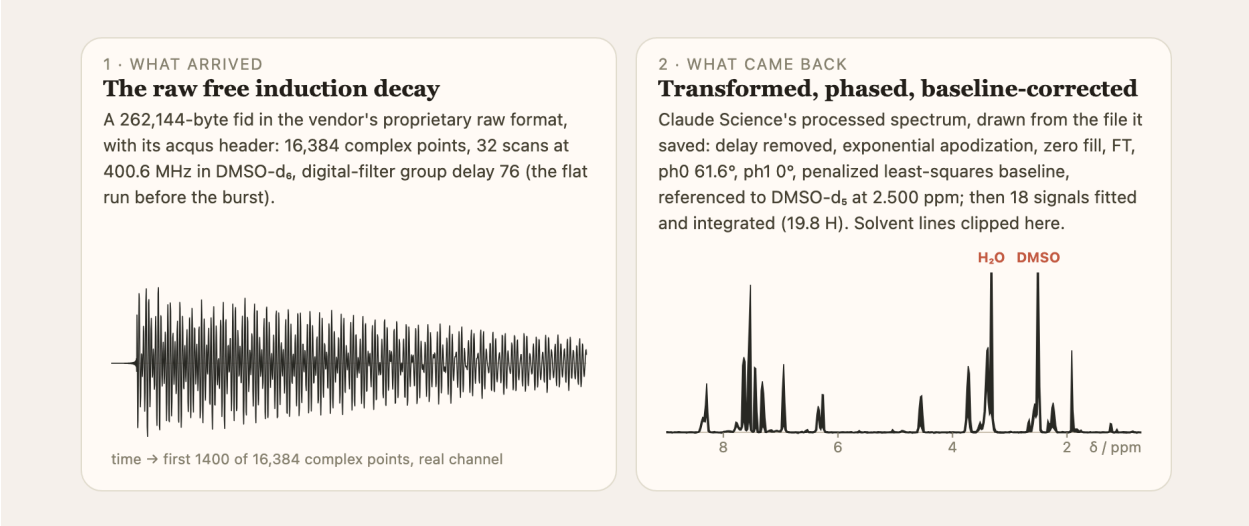

Figure 1. What arrived with Prompt 1 and what came back. Left: the raw FID from the CRO's dataset folder, in the spectrometer vendor's proprietary format. Right: Claude Science's processed spectrum, drawn from the spectrum file it saved.

Processing and peak table

Claude Science took the group delay from the header, applied exponential apodization and zero filling, and phased the spectrum to ph0 = 61.6°, ph1 = 0°; a penalized least-squares baseline was verified flat to 0.15 σ over fourteen signal-free windows, and the axis was referenced to the residual DMSO-d_5 quintet. Lines were fitted as pseudo-Voigt profiles, and a coupling constant was reported only where the reconstructed first-order pattern reproduced every observed line. The whole first answer took 91 tool calls; the resulting table lists 18 signals summing to 19.8 H.

> *Claude Science: "Signals at 8.36, 7.76, 6.34 and 3.52 are broad singlets with no resolvable coupling; linewidths of 12–33 Hz suggest exchangeable protons (NH/OH). A D_2O shake would confirm."*

The CRO had in fact run that experiment three days after the first spectrum, and its FID was supplied next.

The exchange experiment

Told only that the sample was “with D2O & d-TFA”, Claude Science established from the header and from the 1.75 Hz line spacing of the residual DMSO-d_5 (CD_2H) signal that the bulk solvent was still DMSO-d_6 with D_2O and d-TFA added, noted that the run came from a different 400 MHz instrument, and set aside the HDO signal at 5.63 ppm, which carries 64% of all intensity. Its first comparison, delivered 18 min after Prompt 2, carried a figure headline claiming that four exchangeable signals had been removed. Its built-in auditor showed that the headline contradicted Claude Science’s own comparison table and that the table rested on a superseded row set; Claude Science replaced the row-by-row matching with integrals over padded regions common to both spectra and corrected itself: “Two regions clear, not four”. The regions at 8.19 to 8.54 and 6.16 to 6.44 ppm fall from 1.76 to 0.21 H and from 1.83 to 0.02 H, a 7.5 Hz doublet of 0.63 H appears at 5.96 ppm, and the 3.30 to 3.88 ppm region is excluded because it overlaps the residual H_2O signal of the first spectrum (3.32 ppm), so its integral is not comparable between the two spectra. Of the four singlets Claude Science had flagged, two (8.36 and 6.34 ppm) were therefore removed, one (7.76 ppm) was retained at 0.63 H (operator: 0.70 H), and one (3.52 ppm) lies in the excluded region; Claude Science also pointed out that the downfield loss includes a sharp one-proton multiplet and so is not purely exchange. The operator's integrals tell the same story: 1.82 to 0.18 H at 8.4 ppm, no region drawn between 6.1 and 6.45 ppm after exchange where 1.89 H had been, and 0.75 H for the new doublet.

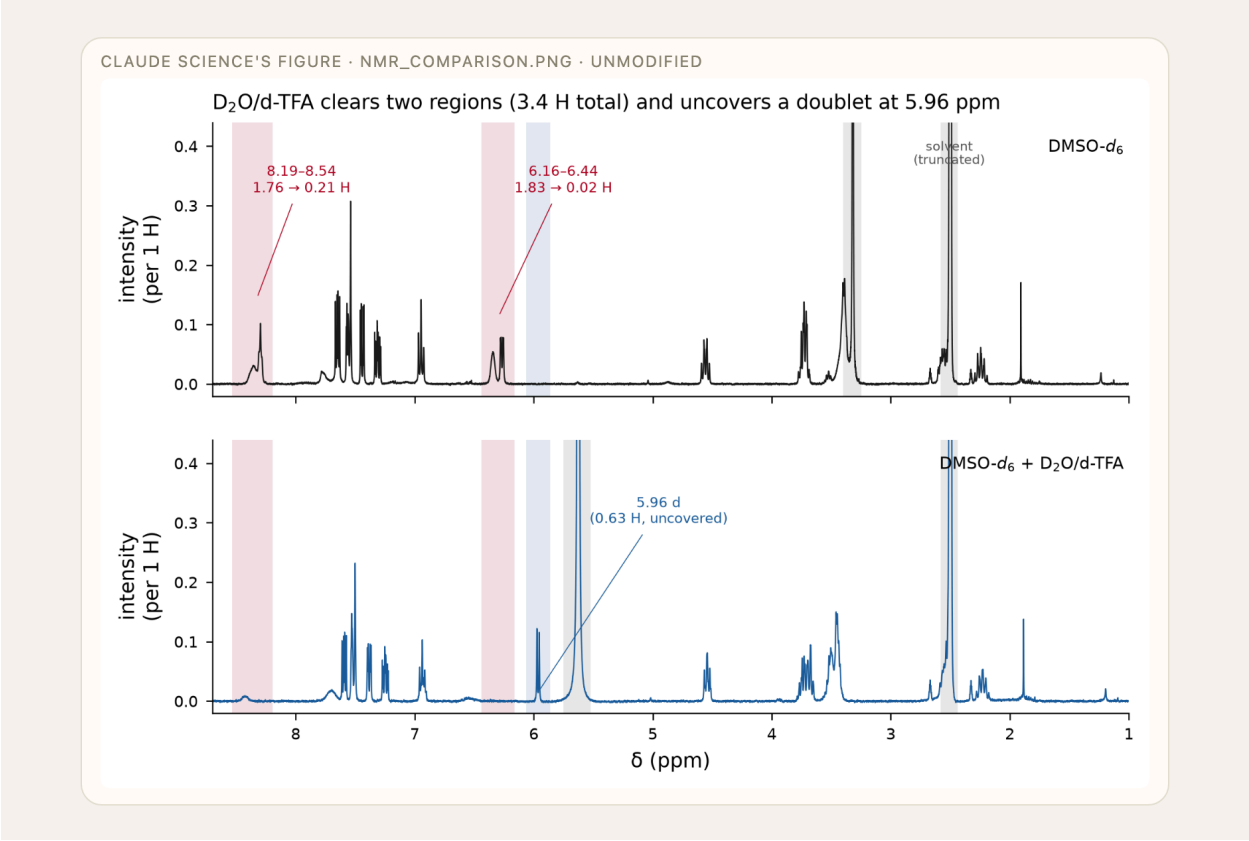

Figure 2. Claude Science's final comparison figure (reproduced unmodified): neat DMSO-d_6 (top) and DMSO-d_6 with D_2O/d-TFA (bottom), normalized per proton, with the two cleared regions and the uncovered doublet marked.

LC-MS : from a vendor binary to chromatograms, spectra and purity

The LC-MS run of the same sample was acquired on July 31 on a single-quadrupole LC-MS system with a photodiode-array (PDA) detector, and the CRO's report page was printed from the vendor's software: a seven-peak PDA table with the main component at 4.34 min and 96.33% area, and ion chromatograms for m/z 505 and 503 peaking at 4.38 min. Claude Science received only the sample.lcd file, in a separate session running alongside the NMR work. It found no parser for the format in its skill catalog or on PyPI, so the job began as one of decoding. Nineteen minutes after the prompt it delivered chromatograms, mass and UV spectra, a purity table and a reusable reader (Figures 3 and 4), and every retention time and ion mass that also appears in the CRO's report agrees with it; the two purity figures rest on different peak sets and are compared below.

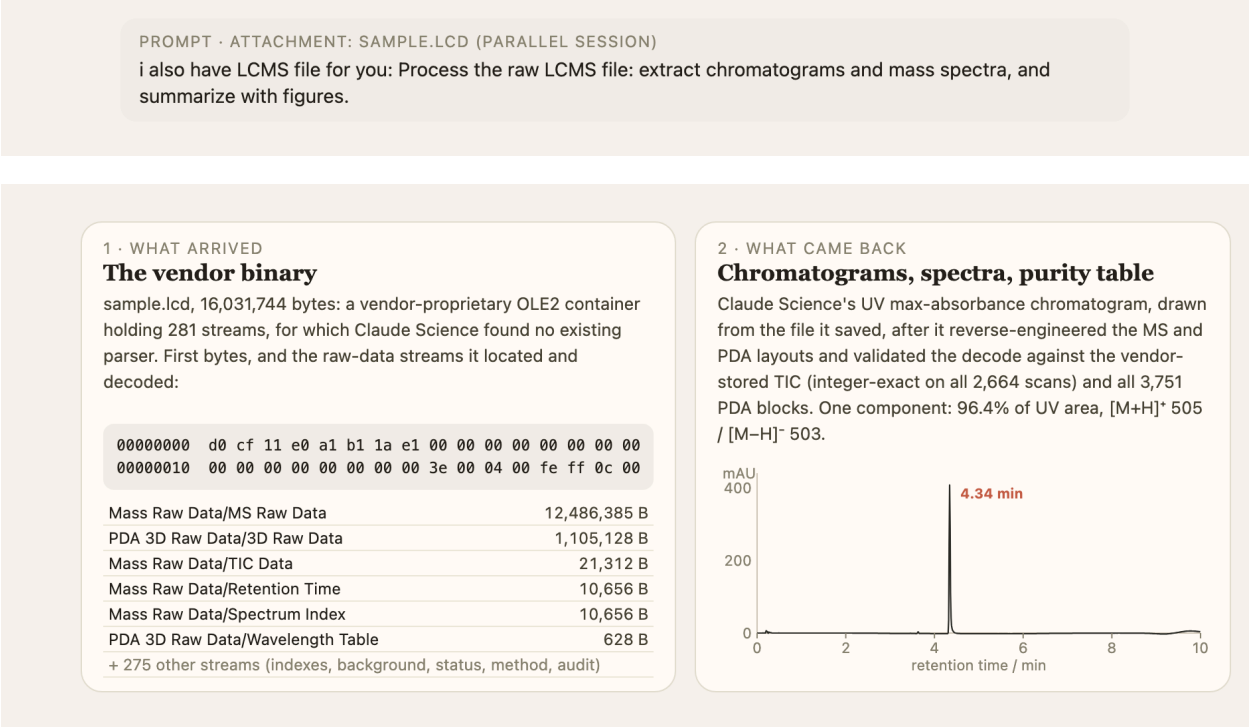

Figure 3. What arrived with the LC-MS prompt and what came back. Left: the first bytes of the lcd container and the raw-data streams Claude Science decoded (stream names and sizes read from the file). Right: Claude Science's UV max-absorbance chromatogram, drawn from the file it saved.

Decoding and result

Claude Science recognized an OLE2 container written by the vendor's proprietary acquisition software, located the MS and PDA raw data among its 281 streams, and deduced their binary layouts. It accepted the decode only after its internal checks passed: the summed ion intensities reproduced the vendor-stored total-ion chromatogram exactly, to the integer, for all 2,664 scans; all 3,751 PDA blocks parsed with exact byte consumption; and the decoded wavelength and m/z axes matched values stored independently elsewhere in the file. The result (60 tool calls): one component at 4.34 min carrying 96.4% of the UV area summed over the three peaks detected from 0 min, void peak included (the CRO's 96.33% is computed over seven peaks from 3.46 min; on that basis Claude Science's own areas give 98.8%, because five of the CRO's minor peaks, 2.1% of area in its table, barely register in the decoded max-absorbance trace); $[M+H]^+$ 505, $[2M+H]^+$ 1009, $[M-H]^-$ 503 and $[M+CH_3COO]^-$ 563, converging on a nominal mass of 504 Da, with ion traces peaking at 4.38 min as in the report; UV maxima at 232 and 290 nm. Claude Science itself stated three limitations: a single quadrupole gives nominal mass only; the intensity-to-mAU scaling was inferred rather than documented; and the vendor's own "Max Plot" stream did not decode, so the max-absorbance chromatogram was recomputed from the validated absorbance matrix.

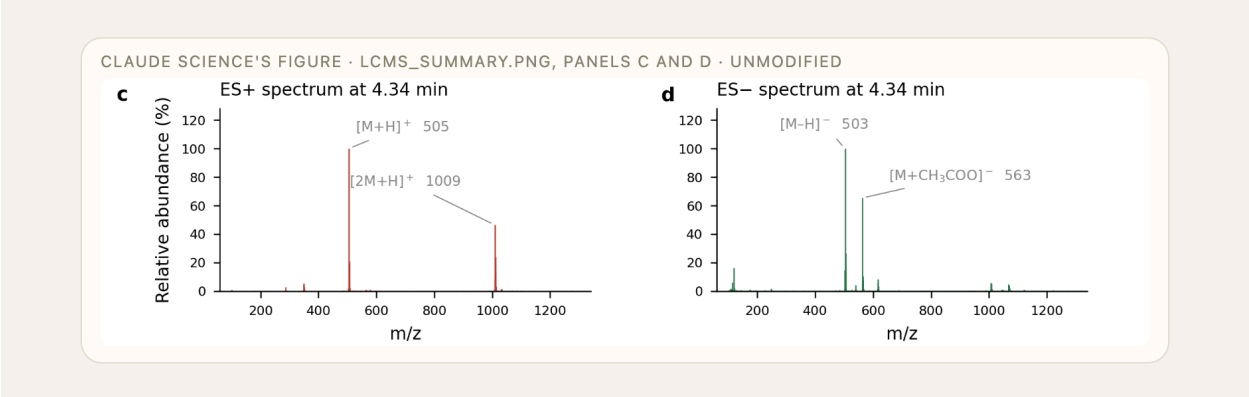

Figure 4. Panels c and d of Claude Science's six-panel LC-MS summary (reproduced unmodified): the ES+ and ES− spectra at 4.34 min with Claude Science's adduct assignments.

Asking about a figure by pointing at it

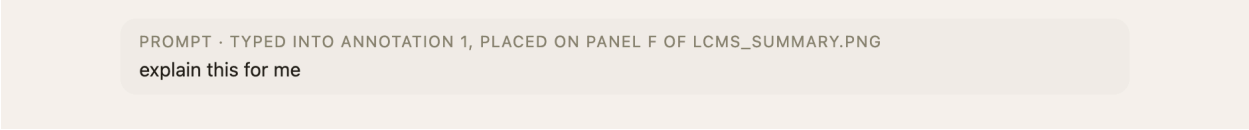

Any figure Claude Science delivers can be annotated in place, and the annotation becomes the prompt. Placing a numbered marker on the extracted-ion panel of the LC-MS summary, with a short request typed into it ("explain this for me"), was enough to get an answer: it came back 29 s and one tool call later. Claude Science located the marker on panel f and then answered from the underlying data rather than from the pixels, re-extracting the two ion traces with that one tool call. It explained what an extracted-ion chromatogram isolates that a total-ion trace does not and cited the peak's 34,000-fold rise over baseline, the 0.98 correlation between two ions measured in opposite polarities, and a 2.3 to 2.5 s delay between the UV cell and the ion source that the plot itself does not show (Figure 5).

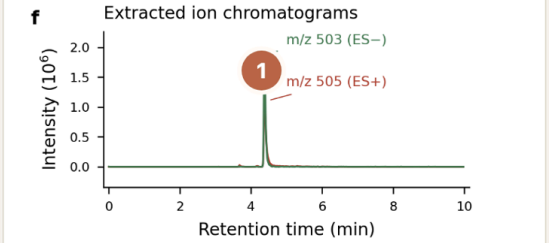

PROMPT · ANNOTATION 1 PLACED ON LCMS_SUMMARY.PNG

explain this for me

Panel f of Claude Science's six-panel LC-MS summary (its panels c and d appear in Figure 4), with the numbered annotation redrawn at the position recorded in the session export (74.0%, 84.9% of the full figure). The request carried no other text.

CLAUDE SCIENCE · REPLY AFTER ONE TOOL CALL (29 S) THAT RE-EXTRACTED THE TWO ION TRACES

[...] Panel f throws all of that away and asks only "how much of *this one mass* is coming off the column right now?" [...]

Each trace sits at a baseline of ~35-56 counts for the whole 10-minute run, then spikes to **1.16 M** (505) and **1.92 M** (503) - about **34,000× baseline** for both. [...] The two traces rise and fall together (correlation **0.98** across the run).

[...] the EIC apexes land at **4.38 min**, while the UV peak in panel a is at **4.34 min** - a **+2.3 to 2.5 s** offset. [...] It's the transfer time through the tubing between the PDA flow cell and the ion source - the detectors are plumbed in series, so the UV sees the band first.

Figure 5. Left: panel f of Claude Science's LC-MS summary with the annotation redrawn at the position recorded in the session export, and the prompt as typed. Right: excerpts of Claude Science's reply; omissions marked [...].