How Doppler geolocation works

An illustrated guide to the Sitrec addon

The heard pitch of an aircraft changes as it passes a stationary microphone. Its recorded frequency curve depends on its motion relative to that microphone. If we have timestamped flight tracks, we can predict what each flight would sound like from many possible listening positions and compare those predictions with the recording.

Sitrec ranks combinations of **listener position, aircraft and recording start time**. A useful match must explain the observed curves and the timing between flyovers. This guide follows the implemented search, from local audio to the labelled result image.

Two kinds of evidence meet in one model

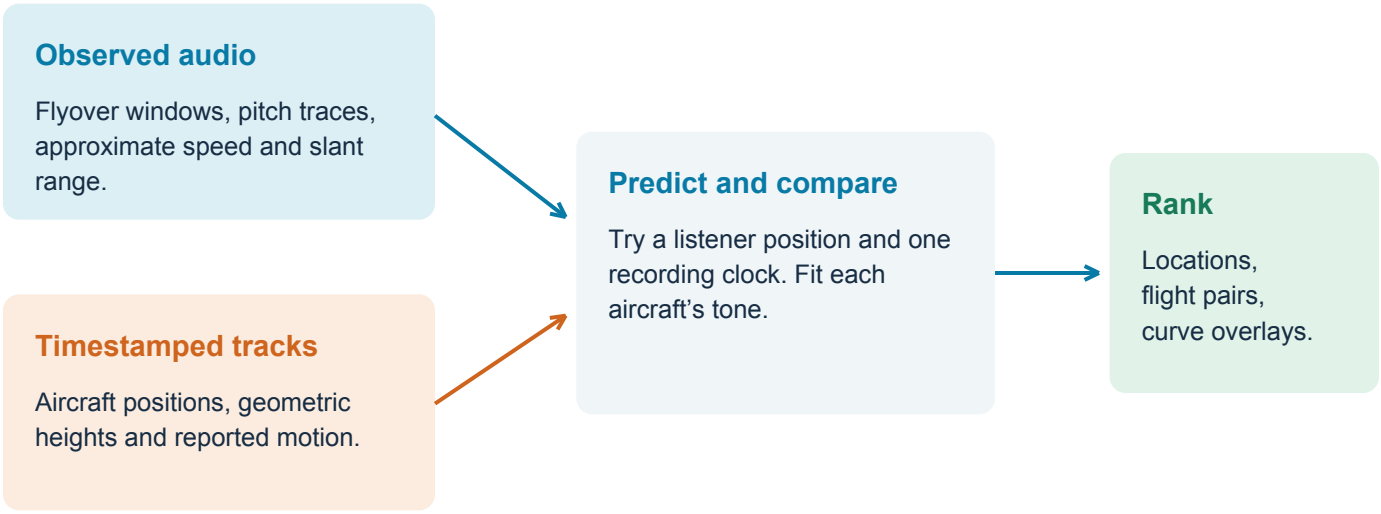


Figure 1. The audio supplies a pattern; the tracks supply possible causes. Matching uses both geometry and timing.

What you need to provide

The recording and its likely date. Use the same continuous recording for both events. For archive discovery, select all parts of the relevant daily flight archive.

Measurements from the spectrogram. Mark each flyover, trace a harmonic and give generous frequency, speed and closest-slant-range bounds. These measurements are currently manual.

A search area or a known clock when available. Continental-US bounds are the default. Unknown-time archive discovery uses two flyovers; a single flyover requires a recording-start UTC.

The aircraft is heard in the past

The position that emitted a sound differs from the aircraft position when that sound arrives.

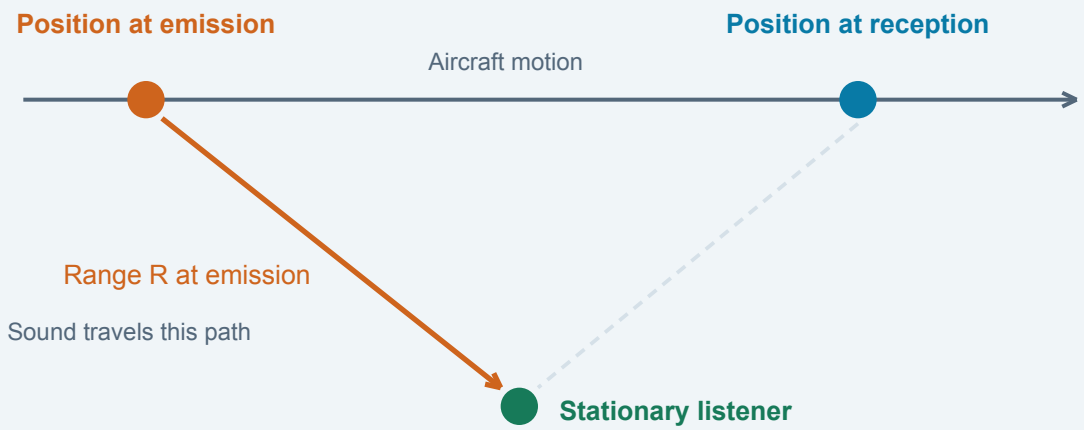


Figure 2. At a 6 km slant range and 343 m/s sound speed, propagation takes about 17.5 seconds. Range changes during the flyover, so one constant time shift is insufficient.

Reception time

$$t_r = t_e + R(t_e) / c$$

Observed frequency

$$f_{observed} = k f_0 c / (c + v_r)$$

v_r is positive when the aircraft is moving away from the listener.

t_e is emission time; **t_r** is reception time; **c** is sound speed. **f_0** is the emitted fundamental frequency and **k** is the selected harmonic number. The radial velocity **v_r** is the component of aircraft velocity along the line from listener to source, evaluated at emission.

Approaching motion makes v_r negative, which raises the heard frequency. Receding motion makes it positive, which lowers the frequency. Sound emitted at geometric closest approach arrives later; the tool reports that delayed arrival as its CPA sound time.

The implementation interpolates 3D positions in Earth-centred coordinates and solves the travel-time equation on each straight segment. For Doppler velocity it prefers reported ground speed, ground course and vertical rate; otherwise it uses position differences. Ground course is not nose heading or airspeed.

Turn the recording into comparable curves

A spectrogram shows frequency vertically and elapsed recording time horizontally.

Decode
4096 Hz mono

→

Window + FFT
Hann window

→

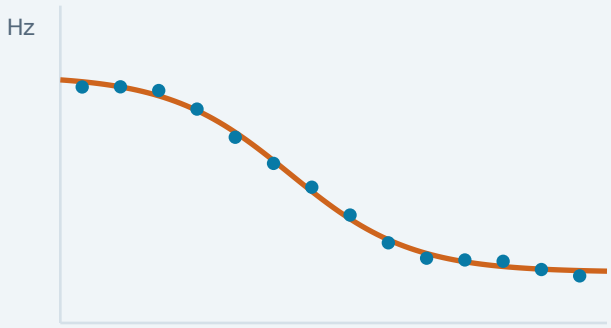
Remove baseline
Persistent tones

→

Normalize
Local contrast

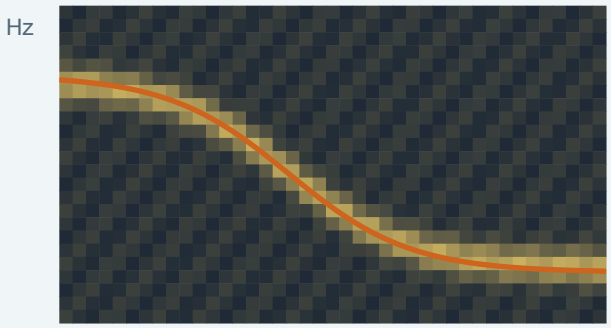
Browser resampling includes an anti-alias filter. Successive FFT windows use Sitrec's shared FFT code. The addon removes the long-term mean at each frequency, then normalizes each time column. Brightness scoring therefore uses processed spectral contrast, not image pixels or calibrated sound pressure.

Trace scoring



Fit to measured dots

Brightness scoring



Sample contrast along the curve

Figure 3. Schematic scoring views. The orange line is the aircraft prediction. Cyan/blue points are manual measurements; the heat field illustrates normalized contrast.

Trace mode fits the allowed emitted frequency to the traced dots and minimizes mean squared frequency error. RMS is reported in Hz. **Brightness mode** tries emitted frequencies and averages spectral contrast along the predicted curve. A higher score is better in both modes.

The combined score gives each event equal weight. With valid traces in every event, an archive brightness search uses those traces to find starting locations quickly, then uses brightness for final refinement and ranking.

The reference used a 4-second FFT window, 1-second hop and 0.25 Hz bin spacing. Bin spacing is not a guarantee of frequency resolution; overlapping 4-second windows also do not provide independent 1-second measurements.

One clock must explain both flyovers

The aircraft tones may differ, but every event must belong to the same recording timeline.

UTC reception time = recording start T_0 + elapsed audio time

Each flyover has its own fitted f_0. All flyovers share the same T_0.

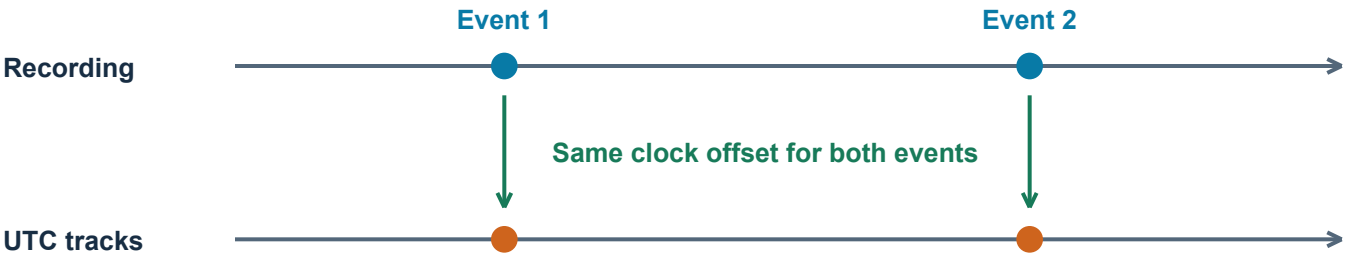
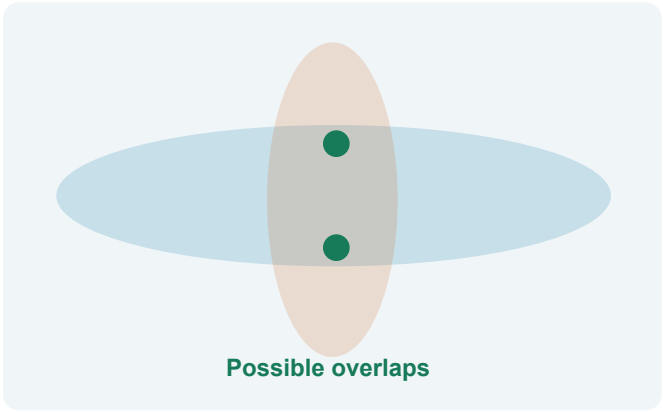


Figure 4. A wrong flight pair cannot be rescued by sliding its two curves independently. Approximate CPA measurements still have their own tolerance.

The delay between flyovers is useful even when the recording-start UTC is unknown. Archive pairing first looks for aircraft that could pass the same broad area with the observed separation. Detailed curve fitting then searches for one listener position and one clock that explain both events.



A single flight can favour multiple positions, including opposite sides of its path. A second flight and the shared clock reduce this ambiguity. Further curve and silence checks can distinguish surviving alternatives.

Figure 5. Schematic spatial ambiguity, not a confidence map or an actual score surface.

Convert an excerpt time to the recording start

Example: if an excerpt begins at audio 05:00 and UTC 20:00:00, the full recording starts at 19:55:00 UTC. Enter that start time, not the excerpt's UTC. Clock uncertainty and CPA measurement tolerance are separate controls.

Search a large archive in stages

The search avoids testing every aircraft at every point across the country.

1. **Read and index the daily archive**

   Stream all split parts. Convert units, keep useful fixes and build a sparse typed-array index. Geometric altitude is preferred.

2. **Pair plausible flyovers**

   Use speed bounds, relative timing and broad spatial proximity. Space/time buckets avoid an all-against-all comparison.

3. **Propose listener positions**

   Start on a 2 km grid near flight-derived anchors, with several height bands. Keep alternative sides and time regions.

4. **Refine every surviving region**

   Use terrain height and continuously vary location and clock. In brightness mode, valid traces can provide a fast starting fit.

5. **Keep feasible curve matches**

   Fit one fundamental per event. Enforce original speed, slant-range and CPA-time bounds. Brightness decides final brightness-mode scores.

6. **Apply optional silence checks**

   Reject a candidate if retained traffic predicts a close approach inside a selected quiet interval, under the chosen audibility assumptions.

7. **Rank and draw score maps**

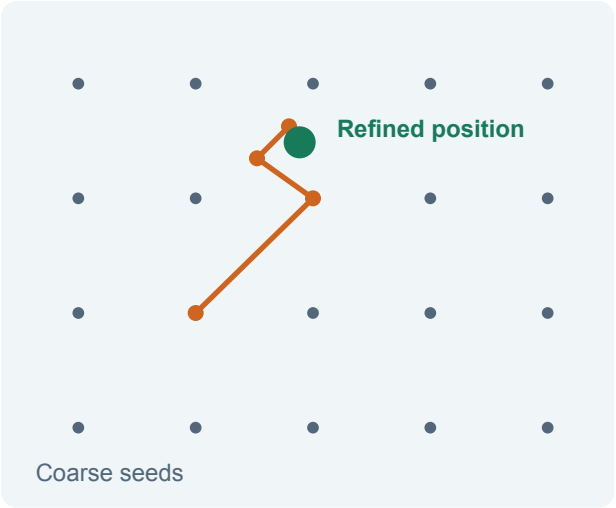
   Keep the best passing candidates and map the leading locations. Result-count and map-count limits do not cut off coarse regions early.

Reference run

79,049 source aircraft produced 42,321 usable tracks, 40,902 pair/time candidates and 7,534 regions for refinement. These counts describe this recording and archive, not fixed search limits.

Refine between grid points and test silence

A coarse grid proposes starting points. It is not the final coordinate resolution.



A bounded simplex optimizer moves the listener and adjusts the shared clock. It can settle between grid points, which is important when a narrow frequency ridge falls between coarse predictions.

Soft penalties steer trials toward the allowed speed, range and timing bounds. The returned solution is the best feasible trial found, even if a penalized trial just outside a bound scores well.

Figure 6. Schematic refinement path. Continuous fitting improves precision, but a local optimum is not a proof of the unique global solution.

Terrain puts the listener and aircraft in the same height system

Flight heights and listener heights must use metres above the WGS84 ellipsoid (HAE). Archive refinement combines terrain height above mean sea level with the EGM96 geoid offset. A negative HAE can still describe a listener at sea level. Pressure-altitude fallback is optional and approximate; current terrain handling clamps below-sea-level terrain to sea level.

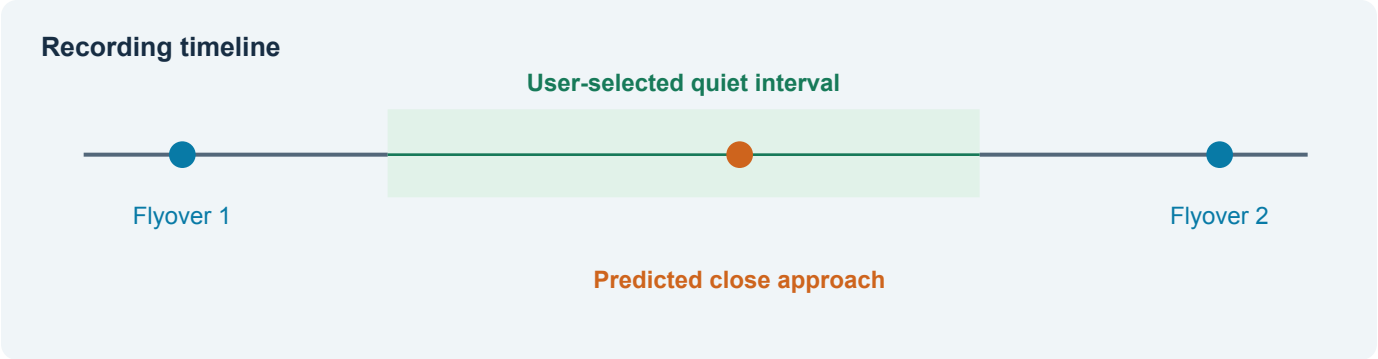


Figure 7. An otherwise good candidate may be rejected if another retained flight should have passed close by during the quiet interval.

Radius zero selects the largest matched CPA slant range. By default, archive silence follows the event aircraft-type filter; excluded types can be included explicitly. This tests predicted close approaches, not guaranteed audible sound. Missing tracks, wind, masking noise or a quiet aircraft can weaken the inference.

Run the analysis in the browser

The main Sitrec startup includes only a small launcher. Administrators load the addon when they open it.

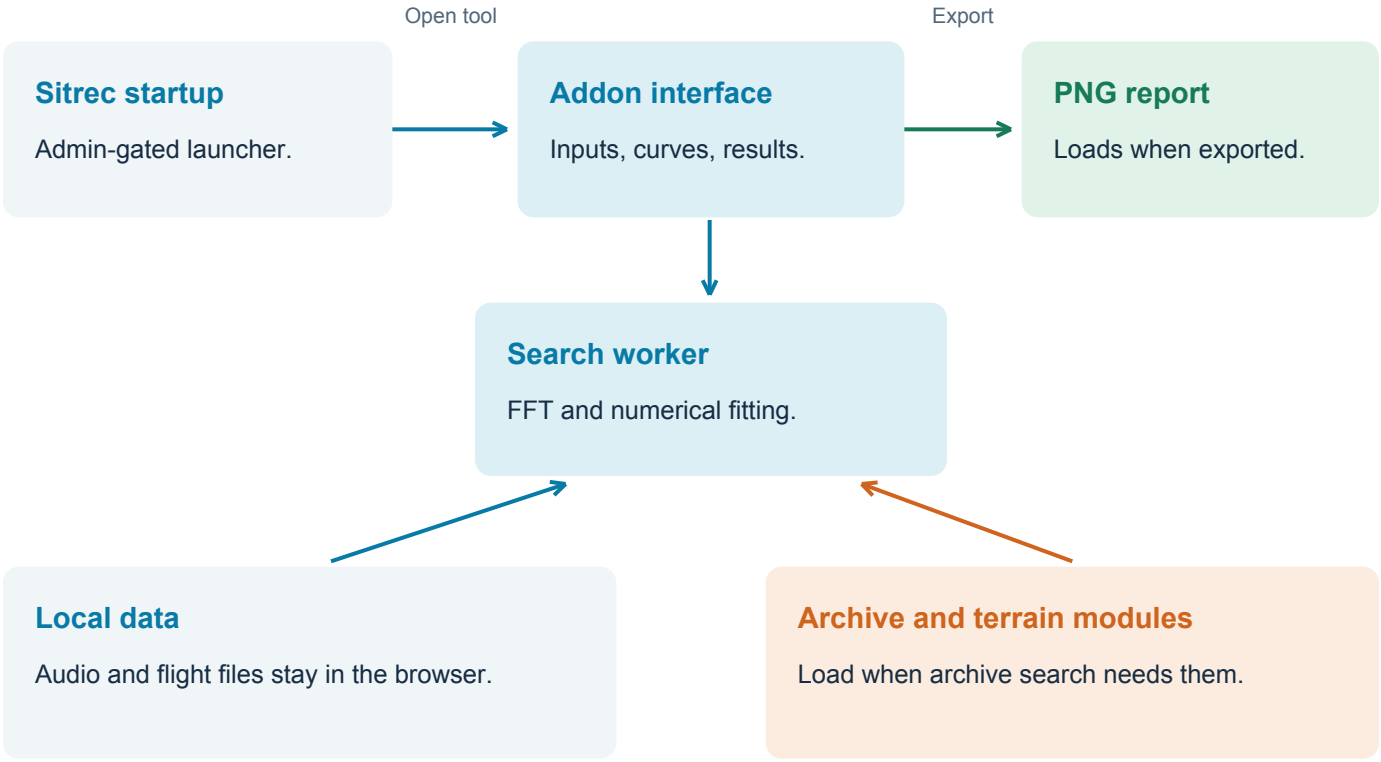


Figure 8. Heavy work runs off the UI thread. Aircraft type classifications and terrain tiles can be fetched over the network; the local recording and archive are not uploaded for scoring.

Why it is faster

Sparse pairing rejects unlikely combinations cheaply. Conservative segment bounds skip distant track segments while preserving the exact closest-approach result. Precomputed spectral interpolation reduces frequency-sweep work. Track, terrain and archive caches avoid repeat reads; independent terrain requests overlap in bounded batches.



Index 94 s | Coarse 214 s | Refine 294 s | Checks/maps 5 s

The complete reference run took **606.6 seconds** with the delivered JavaScript engine and Node file/PNG adapters. This is not a browser benchmark. The full-day index alone used roughly 1 GB; decoded audio and other structures need additional memory.

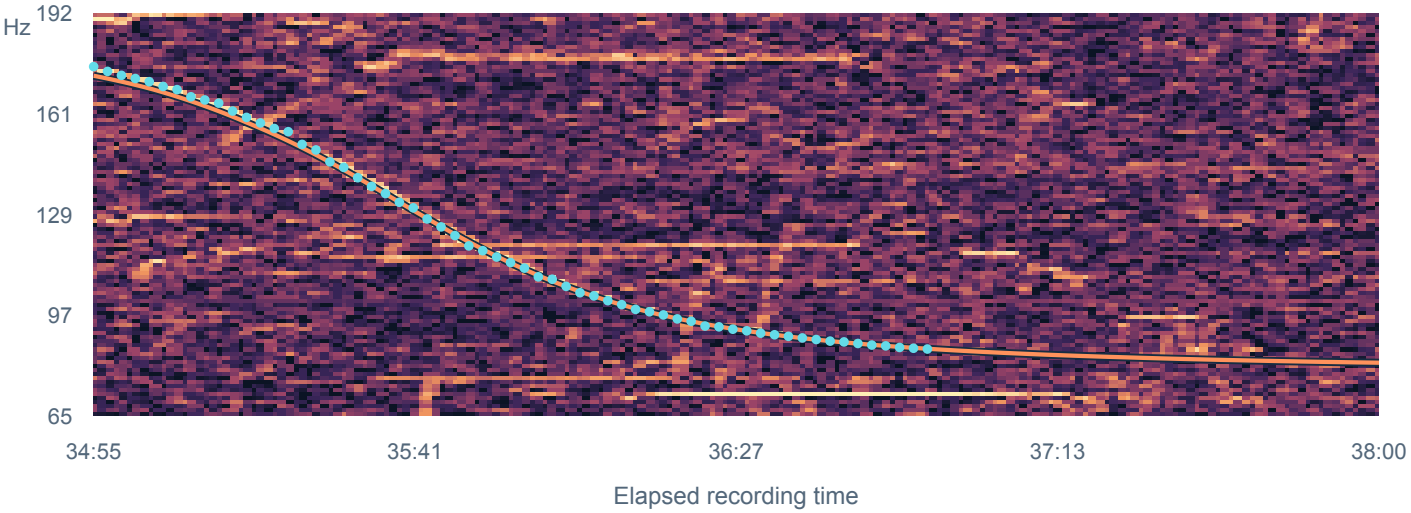
The idle archive index is reused while the tool remains open. Closing the tool releases it. Stopping an active search terminates its worker and requires reindexing. Smaller geographic bounds are the most direct way to reduce work when they are known.

Read the matched aircraft spectrograms

The actual two-flyover result from the supplied recording, with predictions from the committed model.

CST1219 | PC12 | N22NY

Fundamental 112.228 Hz | CPA slant 6134 m | Speed 136.8 m/s



EGE93 | B350 | N93CQ

Fundamental 125.535 Hz | CPA slant 3615 m | Speed 103.9 m/s

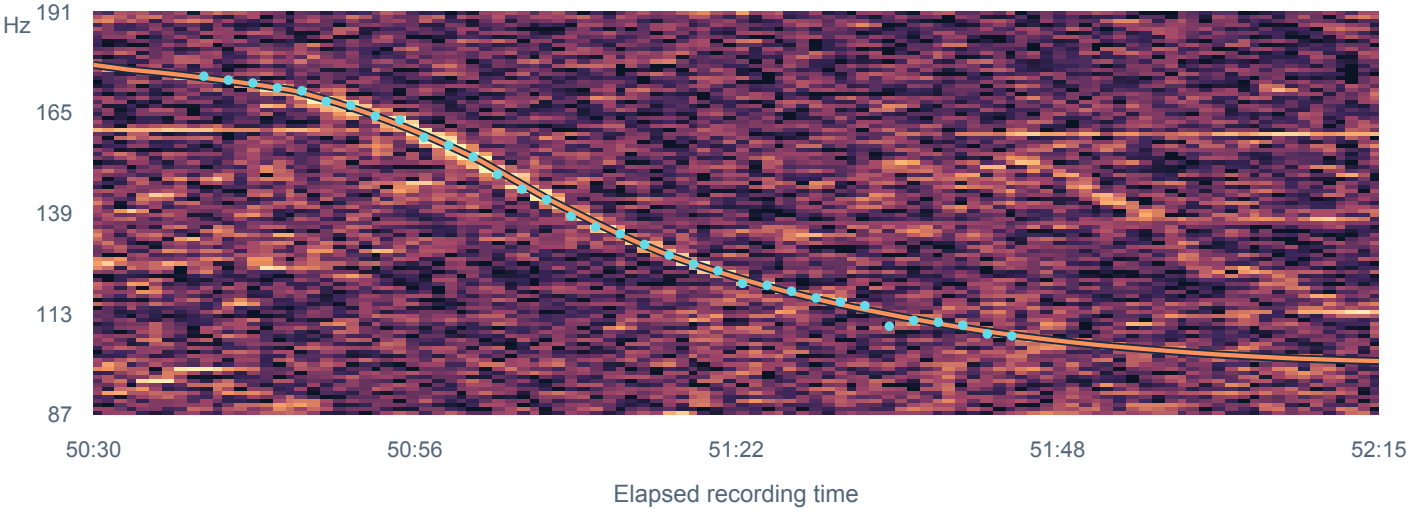


Figure 9. Actual normalized audio contrast with independently fitted aircraft fundamentals and one shared clock. The background is a display of processed contrast, not an audibility threshold.

Top candidate: 40.876770, -73.525160. **Recording start:** 2026-09-08 20:17:29.248 UTC. The two-aircraft score is 2.1487; the next distinct pair scores 2.0612. This margin is not a confidence percentage.

Use the result as evidence

A good curve match is a candidate explanation whose strength depends on the inputs and assumptions.

What the reference run established

A fresh search of the complete supplied 2026 archive ranked the PC12/B350 pair first. The location is **38.1 m from the video's published two-event estimate**. No answer coordinates, aircraft IDs or recording clock were supplied to discovery. Measurement and scoring settings were corrected during development, so this was not a pristine held-out experiment.

The distance measures agreement with the video, not surveyed error to the microphone. The faint single-event example remains unverified. Browser loading of the original audio, saved results, labels and PNG export was checked separately from the timed Node search.

How to inspect a result in Sitrec

Open **Main → Extra Tools → Doppler Geolocation** as an administrator. Load the audio and flight data, choose the scoring mode and mark the events. For unknown time, leave recording start blank and use two flyovers. Review the leading pair and alternatives, then inspect both predicted curves before refining.

The PNG export includes flight identity metadata, fitted frequency, CPA slant, speed and UTC sound-arrival time. Save the analysis to retain measurements, spectral settings, track snippets and candidate results. Reselect local files after reopening; reload the same audio with the saved FFT settings to preserve brightness results.

Assumptions that can change the answer

The model assumes a stationary listener, subsonic flight, a constant emitted fundamental per event and constant sound speed. It does not solve wind advection, refraction, multipath echoes or changing engine tone. Missing fixes, pressure-altitude approximations, a wrong harmonic or a trace following interference can favour the wrong candidate.

Default track gaps above 30 seconds are not interpolated. A finer grid or more decimal places does not establish accuracy. Early spatial and measurement constraints can exclude the true match, and silence requires both adequate track coverage and a defensible audibility assumption.

Sources and implementation snapshot

Code snapshot: local Sitrec commit **bc4ffbfd**. Physics and scoring are in **src/doppler/DopplerModel.js** and **DopplerSearch.js**. Discovery, refinement and terrain are in **DopplerArchiveSearch.js**, **DopplerOptimizer.js** and **DopplerElevation.js** in that directory. The FFT is shared via **src/audio/AudioFFT.js**.

Method demonstration on YouTube | Reference recording on YouTube

Evidence in DopplerGeolocationResults: AcceptanceSummary.json, ReferenceBrowserAnalysis.json, AircraftSpectrogram.png and ProductionBundleAudit.json. Schematics and measured figures are labelled.