

From Warning to Trust

A Multi-Channel Convergence System for Tornado Detection and False-Alarm Reduction, with a Build Guide

Matthew "Bear" Schulz · Independent Researcher · Belmont, CA ORCID 0009-0000-0956-5743 · mattschulz97@protonmail.com

Preprint, not peer reviewed. CC BY 4.0. v3, 2026-08-16. Supersedes v2 (May 22, 2026). Earlier versions remain available and citable. Changes are itemised in Appendix B.

RJ — Remember Joplin. Filed in memory of the 158 people killed by the May 22, 2011 Joplin, Missouri EF5 tornado.

Keywords: tornado warning · false alarm ratio · precision · infrasound · convergence detection · citizen science · Raspberry Shake · noise floor · pilot study · open hardware

Abstract

Tornado warning has two failure modes and only one of them gets attention. A miss is counted; a false alarm is treated as the acceptable cost of caution. But the false-alarm ratio is itself a cause of death — post-event study of the 2011 Joplin EF5 found many residents did not act on a valid warning because prior false alarms had eroded its credibility. A system that improves detection while leaving the false-alarm ratio alone reproduces that problem exactly.

This paper proposes a three-channel convergence detector built from established precursor science — infrasound, lightning rate, and surface pressure shape — in which no single channel raises an alarm and the alert tier is set by how many independent channels agree. It quantifies, for the first time in this work, what that architecture actually buys: with independent channels, requiring two of three to agree cuts the spurious-alarm rate by three to ten times, and with a contaminant that drives all three together it buys **nothing at all**. That result reframes the nine false-positive mitigations from a list of precautions into the mechanism the architecture runs on, since every one of them exists to keep the channels independent.

The author has built and tested nothing. This is a design and a specification, and it is written so that somebody else can build it. Section 6 is a complete build guide with hardware, siting, and a three-step experiment ladder that begins with archived public data and costs nothing. The first experiment is not tornado detection — it is measuring the noise floor, which has never been done for this instrument and which every claim below depends on.

Nothing here is a prevention system. Active-mitigation ideas from the previous version are demoted to open research questions, because no proposed intervention has a validated causal lever and warning-and-sheltering remains the only life-saving intervention available today.

1. The Problem: False Alarms, Not Just Misses

Published operational analyses place NWS tornado-warning probability of detection near 70%, with a false-alarm ratio around 75% [1]. Post-event study of the 2011 Joplin EF5 documented that many residents did not take protective action on a valid warning, in part because prior false alarms had eroded warning credibility — the "cry-wolf" effect [2].

A note on terminology, because the previous version of this paper got it wrong. The 25% figure that complements a 75% false-alarm ratio is **precision**, also called positive predictive value: of all the warnings issued, the fraction that were followed by a tornado. It is *not* specificity. Specificity is $TN/(TN+FP)$, and for tornado warnings the true negatives are every county-hour with no warning and no tornado — an enormous number, which makes specificity already 98–99% before anything is built. Version 2 of this paper set a design target of ">90% specificity," which is satisfied by construction and could never have failed. The target that means something is **>90% precision**, equivalently a false-alarm ratio below 10%, down from 75%.

The premise of this paper is that both axes must improve together, and that precision is the neglected half.

2. Detection Channels, and Why Each One Works

Three channels, each resting on established precursor science. The physical reason each one carries information is given, because a builder who does not know why a channel is there cannot tell when it is lying.

2.1 Infrasound, 0.5-20 Hz

The physics. A rotating vortex radiates acoustic energy at frequencies set by its geometry — larger cores radiate lower. Tornadic core radii put that emission in the infrasound band, below human hearing [3,4].

Why it gives lead time and range. Atmospheric absorption of sound rises steeply with frequency, roughly as the square of it. A 1 kHz sound is gone in a few kilometres; a 5 Hz sound crosses hundreds. This is the same property that lets infrasound arrays monitor nuclear tests at continental range, and it is why this channel can hear a storm that is still far away — the lead time is bought by the physics of absorption, not by cleverness.

What contaminates it. Anything else that rotates or thumps at those frequencies: wind turbine blades above all, plus aircraft, distant explosions, and surf. Turbine contamination is the critical one in tornado alley and is addressed in §4.

2.2 Lightning rate

The physics. A strengthening updraft lofts more graupel and ice, collisions between them separate more charge, and the total flash rate rises. A rapid increase in total flash rate — a *lightning jump* — precedes severe weather intensification.

A correction from the previous version. Version 2 called this channel "Schumann f_1 amplitude." That was wrong for the hardware. Schumann resonances are a *global* mode of the Earth-ionosphere cavity, and a single supercell barely moves the global amplitude. A local ELF magnetometer at a node is not measuring the global mode; it is measuring nearby sferics. Naming the channel for what the instrument actually senses makes it both physically honest and better supported, since lightning-jump behaviour is operational science rather than a novel claim.

Builders should run their own prior-art pass on the lightning-jump literature before writing this channel into a proposal. It is well established, which is a reason to use it and a reason not to claim it.

Free ground truth. The GOES Geostationary Lightning Mapper provides continuous total-lightning data at no cost, which means this channel can be validated against an independent instrument you do not have to own.

2.3 Surface pressure-derivative shape

The physics. A mesocyclone is a low-pressure core; a barometer under one records a dip. The information is in the *shape of the derivative*, not the depth of the dip: a cold-frontal passage produces a slower, one-sided change, while a mesocyclone passage produces a sharper and more symmetric excursion. Absolute pressure varies with weather, altitude and season and carries little; the rate of change and its symmetry carry the signature.

This is an instance of a general rule worth stating for builders: the signal is usually in the rate, not the level. A channel that reads a level competes with every slow drift in the environment. A channel that reads a rate does not.

2.4 What was removed, and why

Two channels present in version 2 are gone.

The ELF 685 Hz channel is withdrawn. Its defining frequency derived from a relation in a companion record that has since been corrected; the relation yields a different number, and the channel has no independent physical basis or established tornado connection.

The 0.47 Hz "beat precursor" channel is withdrawn, for two independent reasons. First, it is not a beat. A beat requires two signals in the *same medium* reaching the *same detector*, where they can mix. Schumann resonance is electromagnetic and lives in the Earth-ionosphere cavity; tornado infrasound is acoustic and lives in the air. An infrasound microbarometer cannot sense an electromagnetic field and an ELF magnetometer cannot sense a pressure wave, and no coupling mechanism between them was proposed. Subtracting one instrument's number from another instrument's number does not produce a physical beat frequency. Second, even taken as arithmetic, it is not stable: 0.47 Hz is the difference between 8.3 Hz and 7.83 Hz, and the tornadic infrasound peak is not fixed at 8.3 Hz — it moves with vortex size across a band, which swings the "beat" through zero.

Recording this removal matters more than the channel did. It was the one element of the design that was original to this work, and it does not survive.

3. The Convergence Engine — and What It Actually Buys

No single channel issues an alarm. The alert tier is set by how many independent channels agree.

Tier	Trigger	Operator meaning
WATCH	2 channels converge	Conditions organising; monitor
WARNING	3 channels converge	Rotation signature likely; notify
IMMINENT	3 channels converge, rotation locked	Act now

Version 2 asserted that this raises specificity. Here is the arithmetic, which the previous versions did not contain.

Model each channel as firing spuriously in a quiet window with probability p . A naive detector alarms if **any** channel fires. The convergence detector alarms on **two of three**. If a contaminant drives all three channels together, two-of-three fires whenever the contaminant does, which is p again.

per-channel p	naive (any fires)	2-of-3, independent	improvement	2-of-3, fully correlated
0.10	0.271	0.028	9.7×	0.100
0.20	0.488	0.104	4.7×	0.200
0.30	0.657	0.216	3.0×	0.300
0.40	0.784	0.352	2.2×	0.400
0.50	0.875	0.500	1.8×	0.500

Read the last two columns together. **Independence buys the entire improvement, and correlation spends all of it.** A wind farm, a solar storm, or an aircraft that fires several channels at once is one vote wearing three coats, and against it the convergence engine is exactly as good as a single sensor.

That reframes §4. The nine mitigations are not precautions layered on top of the convergence engine — **they are what makes it work at all.** Every one of them exists to keep the channels independent, and their real performance metric is not how many contaminants they catch but how far they push the between-channel correlation toward zero.

What the target costs. Take a season with 20 tornadic windows in 1000 and a detection probability of 0.70. Reaching 90% precision requires the joint false-positive rate below about 0.0016, which with independent channels on a two-of-three rule means **each channel must fire spuriously in under about 2.3% of quiet windows.** That is the number the pilot has to measure and the number this design lives or dies on. It is stated here as a requirement, not as an achievement.

4. False-Positive Control

Nine independent mitigations run at all times: wind-turbine geofencing, three-sensor beamforming coherence, full-overtone pattern matching, NWS radar cross-reference, the multi-channel convergence requirement itself, pressure-shape analysis, a solar-storm filter, time-of-day baseline correction, and a duration threshold.

Wind-turbine geofencing is the critical one, because turbine blades produce precisely the rotating-source infrasound signature the first channel hunts for, and tornado alley is full of them.

Per §3, these should be evaluated by the between-channel correlation they remove, not only by the individual false alarms they reject. A mitigation that cuts one channel's spurious rate while leaving it correlated with another channel's has not bought what the architecture needs.

These are design targets to be tested by a pilot, not achieved performance. No field validation has been run.

5. Open Research Questions — Active Mitigation

An honest preface. The items below are questions, not proposals. None has a validated causal lever, because the underlying science — what makes a supercell tornadic, and whether tornadogenesis can be steered at all — is itself unsolved. Component technologies exist; their application to tornado prevention does not. They are listed to invite investigation and better ideas, not to assert solutions. **Warning and sheltering remain the only life-saving intervention available today.**

Question	What would have to be true
Could bleeding a supercell's charge separation, for instance by laser-guided discharge as demonstrated at ~50 m [7], affect mesocyclone maintenance?	That electrical energy is a causal driver of rotation. Not established; tornadoes are shear- and updraft-driven.
Could surface-inflow thermal disruption change the boundary condition that enables rotation?	An energy budget of 100 MW to 1 GW against a system that processes far more, continuously. Unfavourable.
Does any ELF or high-frequency channel correlate with tornadogenesis at all?	No established tornado connection. Speculative, and it belongs here rather than among the detection channels.

The organising idea behind the previous version's mitigation section is worth keeping even though its proposals are not: a Franklin rod does not oppose the lightning's energy, it provides a preferred path, and its leverage is in timing and geometry rather than power. Whether any comparable boundary-condition lever exists for a forming supercell is an open question, and this paper does not claim one does.

6. Build Guide

The author has built nothing. This section exists so that someone else can, without needing to ask him anything.

6.0 Before spending any money — archived data, free, this week

There is an existing global network of low-cost infrasound sensors with open historic data and an API. Raspberry Shake's network carries thousands of devices, many with an infrasound microbarometer, and all historic data is downloadable.

Do this first. Pull archived infrasound from network stations in tornado alley around dated, confirmed tornado events, and ask two questions:

1. Is the 0.5-20 Hz signature actually visible at realistic station-to-storm distances?
2. Across the non-tornado weeks in between, how often does a naive threshold detector fire?

The second question is the important one, because it is the first real measurement of the per-channel spurious rate that §3 shows everything depends on. It costs nothing, needs no hardware, and requires no storm. **If the signature is not visible in real archived data at real distances, stop here** — that is a result, it should be published, and it saves everyone the rest.

6.1 The hardware

Part	Purpose	Notes
Raspberry Shake & Boom, or Boom alone	Infrasound microbarometer, 0.5-20 Hz	The commodity solution. Boards from roughly \$180, complete kits a few hundred. Check current pricing — the figures circulating date from the product's launch. Joins the open network, so your data has somewhere to go.
Raspberry Pi	Host and logger	Required by the above; check the current supported model.
Precision barometer (e.g. a MEMS pressure sensor breakout)	Pressure-derivative channel	Needs resolution fine enough to see the <i>shape</i> of a change, and a stable temperature environment.
Lightning input	Lightning-rate channel	Start with GOES GLM data rather than hardware. It is free, continuous, and independently calibrated. Add a local sferic detector only if the satellite feed proves insufficient.
Enclosure and wind screening	Everything	Infrasound sensors are exquisitely sensitive to local air movement.

Total for one node: a few hundred dollars, not the ~\$1,500 the previous version estimated, because the infrasound sensor no longer has to be built.

6.2 Siting, which matters more than the hardware

Practitioners running these sensors report that careful siting is decisive. A Boom will pick up a door opening, a person walking past, a refrigerator compressor cycling, and passing traffic. Put it on a solid surface, in the quietest available corner, as far as possible from HVAC, appliances, doors and roads. **A badly sited good sensor is worse than no sensor**, because it will produce a spurious rate that looks like a channel.

6.3 The experiment ladder

Step 1 — the noise floor. One sensor, one week, no weather. Record continuously with nothing happening, and measure how much the channel moves when nothing is wrong. Produce a distribution, not a number: the spread is the floor. Everything in §3 is priced against it.

This step is not optional and it comes first. The cost of skipping it is documented elsewhere in the author's work: four separate experiments were run on a channel whose effect turned out to be about a ninth of its own noise, and a single unattended overnight run would have prevented all four. **No detection claim in this paper means anything until this measurement exists.**

The same recording also delivers the time-of-day baseline correction listed in §4, for free.

Step 2 — the contaminant, and this is the elegant one. The dominant false-alarm source is wind turbines, and **you do not need a tornado to test it.** Take a sensor to a wind farm, record the turbine signature, characterise it, then demonstrate that your rejection logic rejects it. Record the same site from several distances so you learn where the contamination stops.

This is worth emphasising because it inverts the usual difficulty. This paper's central claim is about false alarms, which means **the thing it most needs to prove can be tested on an ordinary afternoon at a fixed, safe, known location** — no storm chasing, no season to wait for. Anyone within driving distance of a wind farm can run it.

Step 3 — the pilot. Only now does a multi-node network make sense. Three to five nodes, one tornado season, with every event recorded — hits, misses and false alarms alike — and results published regardless of outcome. The pilot's primary output is not a detection; it is the per-channel spurious rate and the between-channel correlation matrix, because those are what decide whether the architecture can reach the precision target at all.

6.4 What sets each channel's floor, and where each one goes blind

What follows are not measurements. They are the *shape* of what a builder should expect, written down in advance so the calibration week has something to test against instead of merely producing a number with no expectation attached. Where a measured floor contradicts one of these, the measurement wins and the contradiction is worth publishing.

Infrasound — the floor is wind at the sensor. Ground-level infrasound is dominated by wind noise, which is why professional arrays use porous-hose or rosette pipe wind-noise reduction rather than a bare microphone. The competence axis is wind speed at the site, and the uncomfortable consequence can be stated before any data is taken: severe convective conditions arrive with wind, so this channel's floor rises precisely when the channel is most wanted. Range degrades it independently, since the signal falls off with distance while the local wind floor does not. Expect a strong diurnal and seasonal component, and expect wind-noise reduction hardware to be the largest single improvement available.

Lightning rate — the floor is the square root of the background. This is a count channel rather than an amplitude channel, so its noise is Poisson: against a background of N flashes in a window, the fluctuation is of order $\sqrt{N}$, and a jump has to clear several of those to mean anything. The consequence is worth knowing in advance — the detectable jump size grows as the square root of the background rate, so this channel is sharpest in quiet air and dullest inside an already-active storm field. It is also the one channel whose floor can be characterised with no hardware at all, since the satellite record is public.

Pressure shape — the floor has two parts and only one of them is noise. The deterministic part is the semidiurnal atmospheric tide plus synoptic drift. It is subtractable, and removing it is exactly what the time-of-day baseline correction in §4 is for. The random part is dominated by the sensor's own temperature coefficient, and — if the sensor is sited indoors — by building pressurisation, where HVAC cycling and door events produce transients far larger than a mesocyclone signature at range. Reading the derivative's shape rather than the pressure level removes most slow drift, which is why §2.3 specifies it that way.

Why these three channels rather than any three. Each goes blind under a different condition: wind, background flash rate, and thermal drift. That is not a coincidence of the design, it is the design — §3 shows that the convergence architecture buys nothing at all from channels that fail together, and everything from channels that fail apart. It is also why the wind-farm case is easy and the ordinary-thunderstorm case is hard: a turbine lifts one channel's reading, while a thunderstorm lifts all three from a single common cause.

The rule that comes before all of them. Rank each channel by effect divided by floor before running anything. A channel whose expected effect is smaller than its own floor cannot work, and no amount of processing downstream will recover it. That check costs one division and it is the reason step 1 exists.

6.5 What to publish, whatever happens

The noise floor. The turbine signature. The correlation matrix. The false-alarm rate over quiet periods. **A negative result with a floor under it transfers, and a positive result without one does not.** If this design cannot reach 90% precision, saying so with the measurements attached is more useful to the next person than any warning system this paper could have promised.

7. Collaboration Sought

On the detection layer specifically: the Oklahoma State University infrasound group (Elbing), the University of Alabama in Huntsville, the NOAA/NWS Storm Prediction Center in Norman, county emergency-management offices in tornado alley, and the Raspberry Shake citizen-science community, whose existing network is the cheapest available path to a real dataset.

The ask is an open question — can multi-channel convergence cut the false-alarm ratio? — not a finished claim.

References

[1] Brooks HE. Tornado-warning performance in the past and future. *BAMS* 85 (2004). [2] Simmons KM & Sutter D. False alarms, tornado warnings, and tornado casualties. *Weather Climate Soc.* (2009). [3] Elbing BR et al. Measurement and characterization of infrasound from a tornado-producing storm. arXiv:1809.00038 (2018). [4] Bedard AJ. Low-frequency atmospheric acoustic energy associated with vortices. *BAMS* (2005). [5] NWS Operations. Mesoscale analysis techniques. NOAA technical memoranda. [6] Raspberry Shake / OSOP. Raspberry Boom and Shake & Boom infrasound instruments; ShakeNet open data and API. raspberrysshake.org. [7] Houard A et al. Laser-guided lightning. *Nature Photonics* 17, 231 (2023).

A prior-art pass on the lightning-jump literature is outstanding and is flagged in §2.2; the reference list above does not yet include it.

Appendix A — Withdrawal of the Reference Code

Two code files shipped with earlier versions of this record: `HSM_WeatherAlert_v1.py` and `HSM_WeatherAlert_Dashboard.html`. They are withdrawn. They are not included in this version, and they should not be used.

They remain retrievable from the earlier versions of this record, which is why this notice is here rather than absent. Anyone who has already downloaded that crate should read this section before running anything in it.

Why, specifically

The alert-dispatch path names siren and public-alert activation with no interlock. `alert_authorities()` returns recommended actions including "ACTIVATE SIRENS" and "SEND WIRELESS EMERGENCY ALERT," and its own docstring describes connecting to NWS, IPAWS and SMS gateways in production. There is no dry-run default, no operator gate, and no guard against an unattended caller. Wiring that stub would connect an uncalibrated detector — one with no measured noise floor — directly to public alerting infrastructure. **This is the reason the files are withdrawn rather than merely deprecated.**

The convergence engine does not converge. §3 of this paper rests on the claim that no single channel raises an alarm, and computes what that buys. The code's WATCH condition is satisfied by a single channel's score alone, and its ELEVATED tier has no channel-agreement condition at all. The arithmetic in §3 does not describe that program; it describes something closer to the naive detector §3 shows is three to ten times worse.

The experimental-channel fence is declared but not implemented. The alert payload tells its recipient that experimental channels are excluded from tier calculation until validated on fifty events. In the code, the withdrawn 0.47 Hz channel sets a score of 0.3, which is exactly the ELEVATED threshold. The promise and the behaviour disagree.

The lead-time figure was manufactured. The demonstration function contains a hand-written scenario — the spectra, magnetometer values and pressure series at each timestep were typed by the author — and the code prints the gap between that scenario's escalation and the operational NWS average as "ADDITIONAL LEAD TIME: ~47 minutes." A deployment note then calls it a conservative estimate and extrapolates to lives saved per year. It is neither an estimate nor a measurement. A scenario built to escalate escalates.

The scoring functions cannot read a real spectrum. They retrieve amplitudes by exact floating-point dictionary key — 8.3 Hz, 0.94 Hz — which no FFT bin will produce. They work on hand-constructed input and nothing else, which is structural evidence that the detector has never been run against real data.

Further findings, including a sample-rate contradiction of a factor of sixty in the pressure channel and the absence of all nine false-positive mitigations, are itemised in the companion code review deposited with this version.

What this does not affect

The build guide in §6 does not use this code and never did. Reading archived network data, siting a sensor, recording a noise floor and characterising a wind farm are hardware and method. A builder writing their own reader for step zero needs perhaps twenty lines and none of them come from here.

What is worth carrying forward

One idea from the dashboard, which was a visual mockup rather than working software: **it shows the operator why a tier fired, not just that it fired**. A warning that arrives with its evidence is a different object from a warning that arrives as a colour, and in a paper about trust that difference is the whole point. Any future implementation should keep it.

The channel scores were also normalised and kept separate before fusion, which is the right shape and is what would make the between-channel correlation measurement of §3 possible later.

Appendix B — The Convergence Demonstration

A single self-contained HTML file, `convergence_demo.html`, is deposited with this version. **It is not the withdrawn code of Appendix A and shares nothing with it**. It contains no detection logic, no thresholds fitted to data, no network calls, no sensor input and no dispatch path of any kind. It cannot alarm anything, because there is nothing in it to alarm with. Its banner says so before anything else on the page.

What it is for: this paper argues that a reading means nothing until its noise floor is known, and that agreement between channels is only evidence when the channels fail for different reasons. Both are easier to see than to read. In the demonstration each channel is drawn against its own measured floor — the floor is the ground line, and a reading exists only as height above it — and any channel whose floor has not been measured is drawn as having no ground at all and is excluded from the count rather than assumed quiet.

Four scripted scenarios show the cases that matter. A quiet night, which is the state the system occupies almost always and the state whose spread defines the floors. A tornadic supercell, where three channels rise for three unrelated reasons. A wind farm upwind, where one channel rises alone and the architecture correctly declines to alarm. And an ordinary thunderstorm, where all three rise from one common cause, the count says three, the correlation readout says one, and the tier is held rather than issued — which is the §3 result made visible, and the honest hard case rather than the flattering one.

It is a picture of a specification. Nothing in it has been measured.

Appendix C — Changelog from v2

change	reason
"Specificity" replaced by "precision" throughout	The 25% figure is positive predictive value. Specificity is already ~99% for this problem, so the v2 target of >90% specificity could not fail.
Design target restated as >90% precision (FAR < 10%)	The target that can actually be missed.
ELF 685 Hz channel removed	Its frequency came from a relation since corrected in a companion record; no independent basis, no established tornado connection.
0.47 Hz beat-precursor channel removed	Not a beat: electromagnetic and acoustic signals in different media, sensed by different instruments, with no coupling proposed. And 8.3 Hz is not a fixed value, so the difference is not stable.
"Schumann f ₁ amplitude" renamed to lightning rate	Names what the hardware actually senses; a local sensor does not measure a global cavity mode.
Active-mitigation proposals demoted to open questions	No validated causal lever for any of them.
Feasibility scores (8/10, 5/10 ...) removed	Subjective ratings presented as measurements.
"Harmonic" removed from the framework naming	Convergence is real and does the work; the modifier was decoration.
Framework references and internal series citation removed	The record stands alone; the internal citation pointed at a single paper's first version as though it were a series.
§3 convergence arithmetic added	v2 asserted the architecture's benefit; this version computes it, including the case where it buys nothing.
§6 build guide added	The author cannot build this. Somebody else should be able to, without asking him.
Node cost estimate reduced	The infrasound sensor is now a commodity part rather than a build.
§6.4 added — what sets each channel's floor	Gives the calibration week something to test against, and states in advance where each channel goes blind. The three fail under different conditions, which is the independence §3 requires.
Appendix B added — a synthetic convergence demonstration	Shows floor-relative readings and the correlated-contaminant case, neither of which reads well as prose. Inert by construction: no logic, no network, no dispatch.
Reference code withdrawn and removed from the deposit	The dispatch path named siren and public-alert activation with no interlock. Appendix A gives the full reason; the companion code review itemises the rest. The build guide does not use the code.