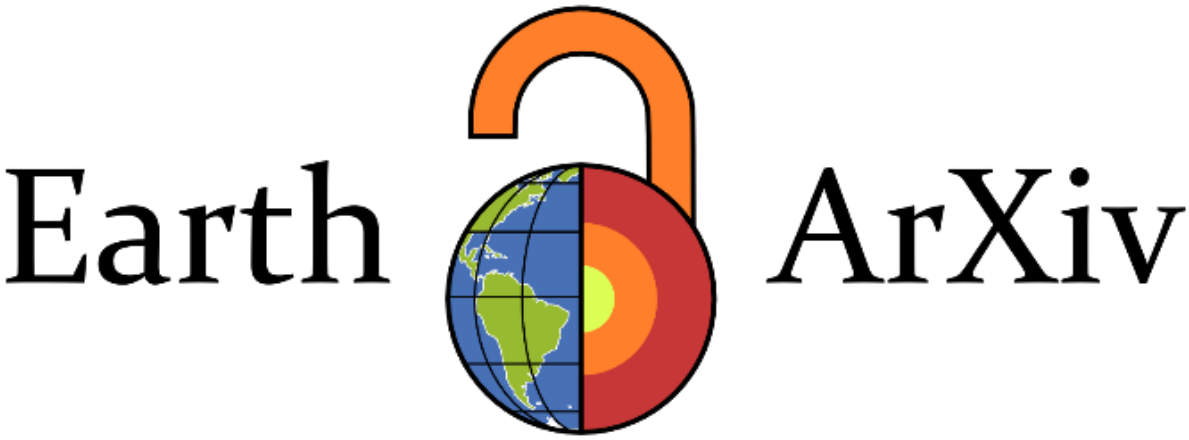

From Public Weather Images to House-Scale Convective Nowcasting: A Latency-Aware, Multi-Source Fusion and Alerting System Built Entirely from Open Data

Antonio Velázquez Bustamante — Facultad de Ingeniería, Universidad de Buenos Aires, Argentina
(avelazquez.ext@fi.uba.ar)

2026-06-11 (v2)

Abstract

Short-range convective warning systems are usually evaluated as gridded meteorological products, but many real decisions are point decisions: whether a specific place will be affected in the next minutes to hours, whether the event is electrically active, and when conditions will improve. This paper presents RADARES, a low-cost operational nowcasting and alerting system built entirely from public data: rendered Argentine weather-radar imagery (SINARAME PNG products), GOES-19 ABI infrared and derived Level-2 products, Geostationary Lightning Mapper (GLM) detections, surface METAR reports, and open numerical-model point forecasts. The system runs 24/7 on consumer hardware, targets a fixed point in northwest Greater Buenos Aires, and pushes phone alerts that answer four questions: when does precipitation arrive, how intense, how long, and when does it improve.

Three design principles distinguish the system. First, **latency and source hierarchy are treated as scientific variables**: the nearest radar (10-15 min latency) triggers urgent alerts while slower context (an official mosaic at 29-37 min; satellite products at 12-60 min) is cached and never blocks the urgent path. Second, **every field recovered from rendered imagery is validated before use**: colorbar inversion carried a silent +5 dBZ bias; ground-control-point georeferencing reached 0.3-1.1 km per-radar accuracy only after associating town location dots rather than text labels. Third, **multi-source physics replaces single-source heuristics** where single-frame discrimination fails: radio-frequency interference indistinguishable from rain bands by geometry alone is removed by physical cross-source vetoes (a $\geq$30 dBZ echo without cloud above, or contradicted by a second radar with a closer view, is physically impossible), validated on 80,444 echoes with 0.0% loss of real-rain echoes.

Retrospective verification shows infrared optical-flow advection beats persistence in MAE at all leads from +10 to +120 min (CSI 0.81 at +10, useful exact-position skill to ~+50 min by FSS), and a method bake-off selects S-PROG over pure advection and ANVIL (CSI 0.39 vs 0.32 at +60 min). Radar nowcasting behaves differently: no fixed method is robust across regimes, and a recent-skill adaptive gate that switches among persistence, advection, and S-PROG never collapses across four contrasting events. Pixel-level convective-initiation detection reproduces the literature's pattern – high detection probability with unusable false alarms – supporting an object-level future path. Operational episodes documented here include a same-day double outage of the radar network and

the model API (the system degraded honestly onto satellite rainfall and a free fallback model source), and an advection-versus-propagation disagreement in which radar echo tracking and satellite field tracking reached opposite, *both correct*, conclusions about system approach – the concrete physical argument for hierarchical fusion without inter-source motion vetoes.

RADARES is not an official warning system. Its contribution is a reproducible, verification-driven architecture for turning imperfect public imagery into actionable point-scale nowcasting – with measured skill, explicit failure modes, negative results reported, and a research roadmap toward object-level and probabilistic convective guidance.

Keywords: precipitation nowcasting; radar image inversion; GOES-19; GLM; S-PROG; pySTEPS; Fractions Skill Score; data fusion; quality control; convective initiation; Argentina; latency-aware warning systems.

1 Introduction

The dominant literature on precipitation nowcasting evaluates the forecast of a field: a radar composite, a satellite image, or a precipitation-rate grid [Pulkkinen2019; Ravuri2021; Zhang2023]. Operational users, however, often ask a narrower and harder question: what happens at one place, soon enough to act? For a household, school, small venue, construction site, or local emergency decision, the useful product is not the most recent radar image. It is an answer to four questions:

1. When will precipitation reach the target?
2. How intense will it be?
3. How long may it last?
4. When is meaningful improvement likely?

The RADARES project was built around this point-decision framing for a fixed site in the Buenos Aires metropolitan region. The available data are attractive but imperfect. Argentina's public radar portal exposes rendered PNG imagery rather than raw radar volumes, and retains only the last ~110 minutes of frames. The closest radar, RMA2 Ezeiza, gives the highest-confidence local signal but has limited range (240 km) and 10-15 min acquisition-plus-publication latency. The official central mosaic sees farther but is lower-confidence and twice as old. GOES-19 full-disk imagery covers Argentina every 10 min but observes cloud tops rather than surface rain. GLM detects lightning directly but is a severity signal, not a rain field. Surface METAR stations report hourly. A free model API supplies point forecasts and a thermodynamic background under a strict daily request budget.

The thesis of this paper is that the useful scientific object is the **fused, latency-aware decision pipeline** rather than any individual nowcasting algorithm. The system does not average sources; it assigns each a role and a confidence semantics (Section 4.5), and it lets disagreement between sources persist as information when the sources measure different physical processes (Section 4.9). The architecture is summarized in Figure 1.

This manuscript contributes:

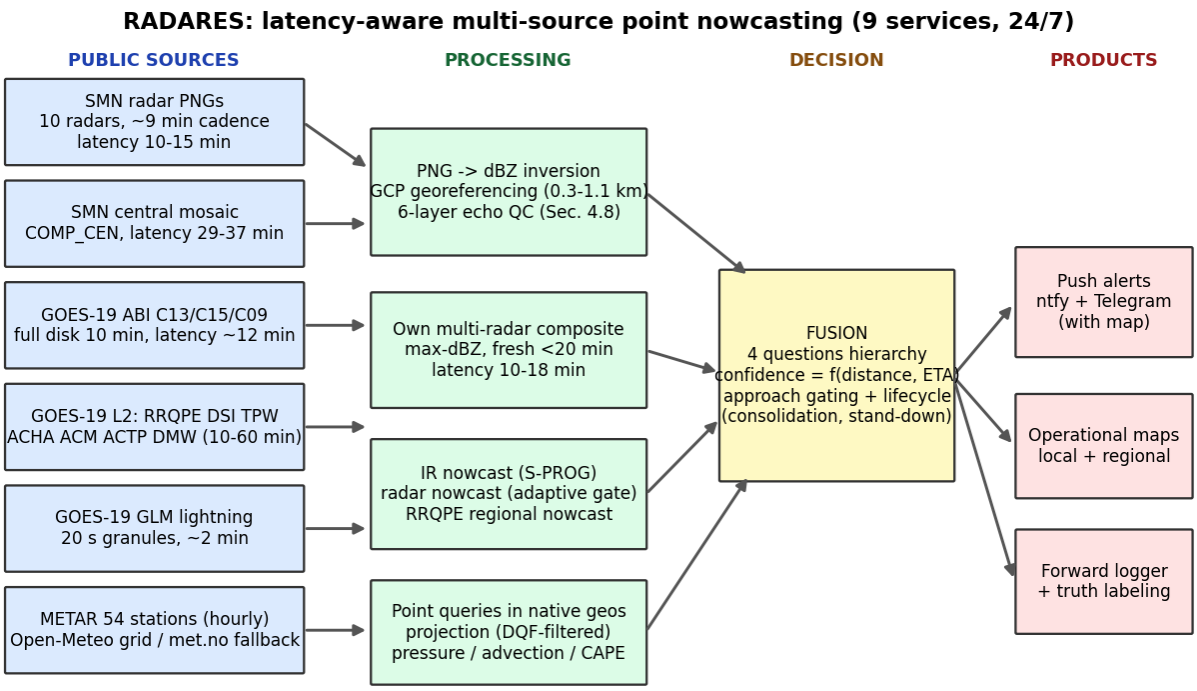

Figure 1: system architecture. Public sources (left, with measured cadence and latency) feed per-source processing including PNG-to-dBZ inversion, georeferencing, six-layer echo quality control, an own multi-radar composite, and three independent nowcast engines; a fusion layer answers the four operational questions under a distance-and-ETA confidence hierarchy with approach gating and a full alert lifecycle; products are push alerts, operational maps, and a forward verification log.

1. **A complete public-data pipeline** for converting rendered radar PNG imagery into georeferenced approximate reflectivity (0.3-1.1 km georeferencing error per radar; a documented +5 dBZ colorbar-sampling bias and its correction).

2. **A latency-aware operational architecture** in which the nearest radar triggers alerts event-driven, while every slower source is cached, fail-open, and monitored, with measured end-to-end improvements (a representative fusion run dropped from 92 s to 3.1 s after the urgent path was isolated).

3. **A six-layer echo quality-control stack** for imagery-derived reflectivity, culminating in physical cross-source vetoes validated on 28 h x 8 radars (80,444 echoes) with zero measured loss of real-rain echoes – including two physically plausible cleaning rules that were tested and *rejected* with data.

4. **A verification-driven method comparison**: S-PROG is the best tested deterministic method for GOES-19 infrared (CSI 0.39 at +60 min vs 0.32 for advection), while radar reflectivity requires regime-adaptive method selection – no fixed method survives all four tested events, and a recent-skill gate never collapses.

5. **An alert-lifecycle design** (approach gating by motion physics and entity size, cross-source

consolidation, and explicit stand-down messages) derived from audited operational episodes, under the principle that *silence after a warning is itself a forecast – an unstated and wrong one.*

6. **A documented advection-versus-propagation case** in which radar core tracking and satellite field tracking disagreed correctly, motivating the design rule that cross-source vetoes apply only to physical impossibilities, never to motion disagreement.

7. **A negative but useful convective-initiation result**: pixel-level cooling-rate detection has signal (POD 0.83 at +10 min) but unusable false alarms (FAR ~1.0), reproducing and reinforcing the object-level consensus.

The intended contribution is deliberately modest and operationally grounded. RADARES is not a replacement for official forecasts or warnings. It is a transparent research prototype for high-value local decision support under data constraints typical of many regions outside dense radar-data ecosystems.

2 Related Work

Classical precipitation nowcasting begins with advection: estimate a motion field from recent observations and extrapolate forward. Lucas-Kanade optical flow [LucasKanade1981] and semi-Lagrangian backward extrapolation [GermannZawadzki2002] remain core ingredients of operational systems. Scale-dependent predictability motivated S-PROG, which decomposes a precipitation field into spectral cascade levels and models each scale with autoregression so that small unpredictable structures decay faster than large ones [Seed2003]. pySTEPS provides open-source implementations of optical flow, extrapolation, S-PROG, STEPS-style ensembles, and verification tools [Pulkkinen2019], and is the algorithmic backbone of this project.

Verification of convective nowcasts requires more than pointwise error. High-resolution fields suffer double-penalty problems: a forecast can place a storm nearly correctly and still receive both a miss and a false alarm at pixel scale. The Fractions Skill Score (FSS) addresses this by comparing neighborhood event fractions across spatial scales [RobertsLean2008]; we use it to separate position skill from structure skill (Section 6.2).

Satellite-based convective initiation (CI) has long used cloud-top cooling, multispectral interest fields, and object tracking to anticipate first radar echoes from growing cumulus [MecikalskiBedka2006; Mecikalski2008]. Operational severe-weather nowcasting systems such as ProbSevere fuse satellite, radar, total lightning, and model environment at object level with machine-learned classifiers [Cintineo2018]. Lightning-jump algorithms relate rapid increases in total flash rate to imminent severe weather with ~20 min lead [Schultz2009]. Overshooting-top detection by IR window and water-vapor channels indicates updraft intensity [Bedka2010]. These works guide our severity layer and our conclusion that pixel-level CI is the wrong operational unit.

Deep-learning nowcasting defines the long-term ceiling: ConvLSTM introduced trainable spatiotemporal recurrence [Shi2015], RainNet matched optical-flow baselines with a convolutional network [Ayzel2020], and DGMR and NowcastNet demonstrated sharp probabilistic radar nowcasts preferred by forecasters [Ravuri2021; Zhang2023]. These require large radar archives and careful probabilistic evaluation; RADARES prioritizes measured deterministic skill and forward data accumulation first.

The data context matters for transferability. SINARAME is Argentina's national C-band radar network [SINARAME]; its public portal exposes rendered CMAX imagery, not volumes. GOES-19 (GOES-East) ABI Band 13 is the 10.3 µm clean longwave window [NOAAStarBand13]; Level-2 derived products (rainfall rate, stability indices, cloud-top height/phase, derived-motion winds, clear-sky mask, total precipitable water) and GLM lightning are distributed openly [NCEIABI; NOAAGLM; NOAAOpenData]. MET Norway's Locationforecast and the Open-Meteo API supply free model point data [MetNo; OpenMeteo]. Many countries publish weather *images* without raw data access; the methods here are designed for exactly that regime.

3 Study Region and Data

3.1 Target and Operational Setting

The target is a fixed point in the northwest sector of Greater Buenos Aires, Argentina (the exact coordinate is anonymized here; the operational system uses it internally). Buenos Aires sits in one of the world's most active deep-convection corridors: mesoscale convective systems formed near the Andes and in northern Argentina propagate east and southeast across the La Plata basin. The dominant alert requirement is not a perfect 2 h field forecast but low-latency detection of near-term threats to one point, with honest confidence.

Figure 2 shows the radar network and the target. The system currently calibrates and ingests 10 radars; the local radar (RMA2 Ezeiza, 49 km from the target) covers the target with high resolution, the others provide a regional curtain out to ~800 km along the climatological approach corridor.

3.2 Data Sources

Table 1 summarizes every source, its measured cadence and latency, and its role. Two latency numbers recur in the design: the local radar PNG appears 10-15 min after scan time, and the official composite 29-37 min after – twice the age. The system's own multi-radar composite (Section 3.3) cuts the regional latency to 10-18 min.

Table 1. Sources, measured timing, and assigned role.

3.3 Own Multi-Radar Composite

Because all individually georeferenced radars share a common 0.02° grid, the system builds its own composite: cell-wise maximum reflectivity over all radars with grids fresher than 20 min, using time windows because radar cadences are not synchronized. The own composite has 10-18 min effective latency versus 29-37 min for the official mosaic, inherits per-radar quality control (Section 4.8), enables cross-radar physical vetoes (Section 4.8, Layer 6), and makes the regional layer independent of the official product – which remains as automatic fallback whenever fewer than two fresh radars are available. The redundancy works in both directions and was exercised by real outages (Section 6.10).

Source	Cadence	Measured latency	Role
RMA2 Ezeiza radar PNG	~9.2 min (phase-locked)	10-15 min	Local trigger: rain now, approaching fronts, intensity
9 further SINARAME radars	8.2-10 min	10-15 min	Own regional composite; transit corridor registration
SMN central mosaic (COMP_CEN)	~10 min	29-37 min	Regional fallback only
GOES-19 ABI C13/C15/C09	10 min (full disk)	~12 min	IR nowcast (S-PROG), clearing, split-window cirrus filter, water-vapor downburst signal
GOES-19 L2: RRQPE, DSI, TPW, ACHA, ACM, ACTP, DMW	10-60 min	12-30 min	Satellite rainfall (third regional source), observed CAPE, precipitable water, cloud-top height (parallax), cloud mask (clearing + echo veto), top phase (hail), measured steering
GOES-19 GLM	20 s granules	~2 min	Electrical severity, lightning jump, truth labels
METAR (54 stations)	hourly	minutes	Surface wind/pressure/temperature; gust-front detection; advection truth
Open-Meteo grid + hourly	20-70 min per ring	seconds	Temperature advection, CAPE/CIN upstream, 48 h outlook (10,000 calls/day budget)
met.no Locationforecast	on failure only	seconds	Free fallback for the model grid (no API key; 20 req/s ToS)
ERA5 (one-time)	static	–	1991-present daily climatology percentiles of the target point

3.4 Satellite Rainfall as a Radar-Independent Regional Source

The RRQPE rainfall-rate retrieval is consumed through the same optical-flow machinery as the radar regional layer, yielding a third regional detection-plus-ETA source that observes precipitation (not cloud) and survives a complete radar-network outage. During the documented national-network outage of 2026-06-11 (Section 6.10), this was for several hours the only regional precipitation source, and it correctly tracked a real rain system approaching from the north.

4 Methods

4.1 Radar PNG to Reflectivity

The extractor identifies the data region and the colorbar, maps RGB colors to nominal dBZ levels, and assigns each pixel to the closest color class. The most consequential calibration lesson was that sampling the colorbar at band *edges* introduced a uniform bias of roughly **+5 dBZ** – one full color class – which shifted every absolute metric while emitting no error. The corrected extractor anchors sampling at band centers and validates that 30 dBZ maps to yellow and 25 to green. The same correction was applied to all four extractor variants (three radar PNG formats and the mosaic). Reflectivity recovered this way is quantized to 5 dBZ classes; all downstream logic is designed around thresholds ($\geq$30 dBZ for alerting, $\geq$35 for truth labels) rather than continuous values.

4.2 Radar Georeferencing

Each radar PNG is georeferenced by combining the PPI disk geometry (Hough-detected center and radius; the 1341x1336 px layout with a 520 px / 240 km disk is shared across the network), the radar site coordinates, and an affine fit in a local azimuthal-equidistant frame against ground-control points (GCPs): towns visible in the imagery. The critical practical detail is to associate each town with its **location dot, not its text label** – text association produced a systematic ~5 km offset with deceptively low residuals. Compact-dot association (circularity and area filters) reduced calibration error to 0.31-1.06 km mean (Table 2). Validation is non-circular: independent markers at known coordinates are rendered over the data and inspected, rather than trusting fit residuals alone. A further lesson: GCPs rejected as outliers by a first contaminated fit can be *valid* – they must be re-projected against the final affine before being discarded.

Table 2. Per-radar georeferencing error (mean over GCPs, km).

Radar	Mean error	Max	GCPs
RMA2 Ezeiza	0.53	0.80	7
RMA1 Cordoba	0.82	1.85	9
RMA6 Mar del Plata	0.47	0.90	7
RMA8 Mercedes	0.31	0.59	6
RMA14 Bolivar	0.88	2.22	6
RMA17 Alejandro Roca	1.06	2.53	10
RMA10 Bahia Blanca	0.72	1.64	7
PAR Parana (INTA)	0.76	1.19	7
RMA18 Santa Isabel	0.49	0.87	11
PER Pergamino (INTA)	0.70	1.35	7

The official mosaic, lacking PPI geometry and town dots, is georeferenced by an affine fit on province-capital markers: 8.6 km mean error (22 km at corners) – adequate for "a system approaches from the west", not for point rainfall, which is exactly its assigned role.

4.3 Infrared Nowcasting

The infrared nowcast uses recent GOES-19 Band 13 frames regridded from geostationary scan angles (x/y multiplied by perspective-point height, transformed with PROJ) to a 0.02° grid. For advection experiments, a dense Lucas-Kanade field is estimated and the last field extrapolated semi-Lagrangianly. S-PROG expects precipitation-like fields, so brightness temperature is transformed:

```
I = max(TB_REF - Tb, 0), with TB_REF = -20 deg C
```

S-PROG runs on `I` and the output is transformed back. The transformation makes warm cloud-free background transparent to the convective-threshold decision; its known cost is that clear-sky Tb saturates at -20 °C after inversion, which does not affect storm thresholds.

4.4 Radar Nowcasting with a Recent-Skill Adaptive Gate

Radar nowcasting runs on dBZ with a $\geq$35 dBZ event threshold and three candidates: persistence, Lucas-Kanade advection, and S-PROG. Unlike the infrared case, no fixed method was consistently

best (Section 6.5). The implemented solution is an adaptive gate: evaluate which candidate best predicted the most recent verifiable segment, and use it for the current forecast. A smoothed flow-coherence diagnostic explains most regime differences – clean advective fields support advection/S-PROG; noisy convective or hail cases favor persistence (where S-PROG collapses to CSI 0.01-0.06).

The urgent alert fusion deliberately consumes *direct fresh radar grids*, not the nowcast product: a continuous-sequence nowcast can be stale after gaps, while the direct grid answers "what is close to the target now?" unconditionally.

4.5 Distance-Time Confidence and the Regional Layer

Sources contribute by distance band, not by identity: anything within 240 km is the local radar's jurisdiction; regional sources only report beyond it (no duplication). Alert confidence combines two independent axes – detection precision (distance: near is sharp) and extrapolation reliability (ETA: systems hours away can dissipate or rain out en route). The confidence is the weaker link of the two:

Distance / ETA regime	Meaning	Priority
<60 km and near-term	Local precision	Urgent
60-180 km and ETA <4 h	High confidence	Urgent
180-240 km	Medium	Normal
>240 km or ETA >8 h	Low	Silent/contextual

Regional ETAs are smoothed by a 3-cycle median with an object-change reset (a >150 km jump in distance resets the history, preventing the median from blending two different systems into a physically impossible arrival speed).

4.6 GLM Electrical Severity

GLM never delays rainfall alerts; it grades them. Flashes within 30 km of the target elevate priority; the flash rate of the nearest cluster provides a continuous severity scale; a Schultz-style lightning jump (2 σ over the recent rate) marks probable intensification ~20 min ahead [Schultz2009]. Truth labeling distinguishes "no lightning" from "no coverage" via window-completeness flags.

4.7 Latency-Aware Alert Chronology

The operational chronology is engineered so that nothing slower than the local radar can delay an urgent alert:

```
RMA2 PNG appears
  -> RMA2 georeferenced FIRST (before all context radars)
  -> alert_loop is event-driven by the new local grid (plus a 10 min watchdog
     so a dead radar cannot silence the other signals)
  -> fusion reads the fresh local grid plus the LATEST CACHED composite, IR,
     satellite-rain and GLM products
  -> the push is sent with the latest available map; the new map renders after
  -> regional refresh, truth labeling, pruning run in separate services
```

After isolating the urgent path, a representative fusion run dropped from ~92 s to ~3.1 s. The general principle: *adding good contextual data must never delay the only source required for the urgent decision.*

4.8 Echo Quality Control and Interference Removal

Reflectivity recovered from rendered public imagery inherits every artifact of the source: speckle, anomalous propagation, biological scatter, and – most damaging – radio-frequency interference (RFI) from radio local-area networks (RLAN), which appears in C-band radars as narrow radial wedges of moderate reflectivity (Figure 5a). Because alert logic consumes ≥30 dBZ echoes, any artifact crossing that threshold is a candidate false front. The system applies six cleaning layers, each validated against a real-precipitation case under a zero-storm-loss requirement.

Layer 1: morphological opening. Isolated ≥30 dBZ pixels are removed by binary opening. On operational grids, noise-only frames contained 4-32 scattered pixels reduced to zero, while real systems (100-850 connected pixels) survived intact; the cut sits in the only natural gap of the blob-size distribution.

Layer 2: physical velocity cap. Lucas-Kanade flow over high-resolution reflectivity produces spurious vectors (measured up to 573 km/h; 3.65% of echoes exceeded 120 km/h). Echoes whose apparent velocity exceeds a physical cap are excluded from approach/ETA estimation. The cap is dynamic when fresh infrared data exist: cloud-top motion is the physical upper bound of cell motion (a cell cannot outrun its own top), with a 20% margin; otherwise a fixed 120 km/h bound applies.

Layer 3: static interference masks by persistence. When the Parana (PAR) radar rejoined the network it exhibited fixed RLAN wedges: 935 grid cells held ≥30 dBZ in ≥80% of 15 consecutive clear-sky frames, organized in 1-5° wedges spanning 27-240 km in range. A static mask built from clear-sky persistence removes them; validated against a storm superimposed on the wedges, the mask removed 992 in-mask cells while 100% of storm echoes outside the mask survived. A non-obvious lesson: **interference masks do not transfer between products.** The official mosaic composites the affected radar unfiltered, and translating the radar-grid mask through geographic coordinates fails (the mosaic's independent georeferencing differs by 8-20 km locally; dilation does not converge). Each product needs its own pixel-space persistence mask, self-aligned by construction. For wedges that wander in azimuth within fixed sectors, per-azimuth-sector persistence masks work where per-cell persistence fails (the wedge flickers cell-wise but inhabits a stable sector).

Layer 4: geometric anti-spoke filter for intermittent RFI. Static masks cannot cover interference that changes azimuth as emitters switch on and off. Per frame, echoes near an affected radar are grouped into contiguous 1° azimuth bins; a group no wider than 4° is a spoke if its radial span is ≥25 km **and** ≥4x its tangential width. Real cells are roughly isotropic and frontal bands are wide in azimuth, while RLAN spokes point at the radar with ~6:1 elongation. An earlier fixed-minimum-length criterion captured nothing (intermittent interference arrives as 30-50 km segments): *relative elongation, not absolute length, is the discriminating feature.* On the contaminated frame the filter removed 60% of remaining spurious echoes; on storm frames it removed none (868/868 grid cells preserved).

Layer 5: alert-level temporal persistence. Independently of field cleaning, the decision layer refuses to alert on the first appearance of an approaching echo lacking support in the previous frame (dilated for expected motion). Ephemeral anomalous-propagation echoes fail this test; a real front confirms one cycle later, costing one radar cycle of lead for the suppression of single-frame false

fronts.

Layer 6: physical cross-source vetoes in the composite. A later interference episode exposed the limit of single-frame geometry: RLAN arrived as a ~9° band that the anti-spoke width exemption deliberately spares, and an attempted sub-windowing of wide groups was rejected by the regression test – real precipitation contains radially aligned groups with 5-16:1 elongations, and the modified filter deleted 97% of a real-storm frame. **There is no single-frame geometric discriminator between an RLAN band and a radially aligned rain band.** The robust solution is physical, applied where sources overlap (the own composite): (i) a **cloud-mask veto** – a ≥30 dBZ echo requires precipitating cloud, so an echo where the ACM clear-sky mask reports clear within a dynamic radius (20 km cloud parallax + 1 km per minute of mask age) is physically impossible and removed; and (ii) a **multi-radar veto** – moderate rain in a cell covered by a second fresh radar at equal-or-shorter range must be seen by both, so a contradicted echo is removed, with a temporal displacement buffer (scans differ by up to 20 min and echoes move; the contradicting radar must be dry over the entire radius reachable at 1.5 km/min – without the buffer the veto deleted 2.2% of a real-storm composite, with it 0.2%). The local alerting radar is never vetoed. Validated against 28 h of cloud-mask history across 8 radars (350 frames, **80,444 echoes**): **0.0% false vetoes on radars with real rain** (0 of 2,559 on the alerting radar), while flagging 37.6% of the most affected radar's echoes as physically impossible – and incidentally discovering a morning anomalous-propagation episode (1,342 ground echoes under clear sky) on a radar not previously known to be affected.

Two physically plausible rules were tested and **rejected with data**. The infrared-consistency rule ("echo without cold cloud above is clutter") fails in winter: measured clear-sky brightness temperature (~+9 °C) sits within a few degrees of shallow stratiform tops, so no Tb threshold separates them – a warm threshold removes nothing, a cold one deletes real rain (the binary cloud *mask* veto of Layer 6 succeeds precisely because it does not rely on a temperature threshold). Static clutter-floor subtraction is pointless on the public product: across 892 consecutive frames only 2-4 pixels showed echo in more than half – fixed clutter is already suppressed upstream, and the residual contamination is intermittent with no stable floor.

The governing principle: every cleaning layer must demonstrate zero storm-echo loss on a real precipitation case before deployment, because a QC stage that silently deletes real echoes is strictly worse than the artifact it removes.

4.9 Approach Gating, Alert Lifecycle, and Advection-versus-Propagation Disagreement

Approach gating. Early regional alerting used the weakest possible criterion: any echo with a positive velocity component toward the target (>5 km/h) counted as approaching, and the nearest such echo defined the ETA. An audited overnight episode showed the failure mode: a real rain system 400 km away advanced 150 km in 3 h (correct low-confidence alerts), stalled at ~250 km and dissolved at dawn – while a separate frame announced a "front" with a 50 h ETA derived from scattered echoes drifting tangentially. The production criterion now requires (i) an approach component ≥10 km/h, (ii) a motion heading within 60° of the target bearing (cosine ≥ 0.5) – removing grazing systems whose projection is technically positive but meteorologically irrelevant; (iii) a minimum **entity size scaled by distance** (beyond 240 km, ≥60 connected approach cells, ~250 km^2: small isolated echoes do not survive a multi-hundred-kilometre journey, and a genuinely growing system qualifies one or two cycles later); and (iv) an ETA ceiling (18 h) beyond which extrapolation is noise. Map captions apply the same gate as the push channel, eliminating map/alert

contradictions.

Alert lifecycle: consolidation and stand-down. Three regional detectors (radar composite, infrared, satellite rain rate) originally alerted independently – three pushes in 45 minutes for one system. Detections are now consolidated: if another source holds an active announcement with an arrival estimate within 90 minutes, the new detection registers silently – one system, one voice. Symmetrically, a system that its announcing source stops seeing as approaching for >45 minutes (a several-frame grace period against flicker) triggers one explicit stand-down – "the announced regional rain is no longer observed approaching" – and clears deduplication state so a reappearance is a new episode. The stand-down closes a real gap: the user woke to sunshine holding three unresolved warnings, and *silence after a warning is itself a forecast – an unstated and wrong one.* Consolidation also applies at closure (no duplicate stand-downs while a sibling source still tracks).

Advection versus propagation: why satellite motion is never slaved to radar motion. A daytime case exposed an apparent contradiction that is in fact the physical justification for multi-source fusion (Figure 7). The radar echo tracker showed the cores of a large convective complex moving east (away) and correctly announced nothing; simultaneously the satellite rain-rate and infrared cloud-top fields both measured a southward component of ~56 km/h toward the target, and surface pressure at the target fell 3 hPa in 3 h – independent in-situ confirmation. **Both motion estimates were correct**: individual cores advect with the mid-level steering flow, while the system as a whole *propagates* by discrete new-cell growth on its moisture-inflow flank – a standard property of organized convection. A core tracker cannot see propagation; a field tracker can. The design consequence: the system deliberately does **not** suppress satellite approach signals when the radar tracker disagrees. Each source announces with its own motion vector; the operational map renders each announcing source's object (dashed ring, source-colored vector and label) alongside the radar arrows; the confidence hierarchy of Section 4.5 arbitrates. Cross-source vetoes (Layer 6) are reserved for physical impossibilities, never for motion disagreements – which are often the signature of propagation, and therefore information.

Display semantics. Motion arrows on operational maps encode *visual ETA*: each arrow's length is the theoretical 60-minute displacement, so time-to-arrival can be read geometrically. Arrows are therefore drawn only for echoes satisfying the approach gate – everything else is scenery and was observed to mislead – with one exception: within 30 km of the target the arrow always draws, because there the heading itself is the decision-relevant information.

4.10 Source Health, Fallback Chains, and Operational Resilience

A multi-source system is only as honest as its knowledge of which sources are alive. Three mechanisms enforce this:

Degradation doctrine. Every signal is generated by the most precise source available, verified across sources where pertinent, and replaced by an explicit fallback chain when the primary fails – or *not given at all* ("UNKNOWN: radar data 23 min old" is a valid and common output). No source failure is allowed to block the alert path (every consumer is fail-open with per-source staleness thresholds), and stale data is never presented as current.

Source-health monitor. A background monitor checks the freshness of all 17 feeds against their natural cadence; a stalled feed triggers a push stating the impact and the active fallback (re-notified every 12 h while down; recovery is also pushed), and every alert carries a "degraded sources" line while any feed is down. The radar-network health is measured on the *second*-freshest radar, since the

composite needs two. Rate-limited APIs get earlier warning: a quota rejection (HTTP 429) notifies immediately (it is a budget problem, not weather), network failures on the second consecutive failure.

Forensic continuity. Service logs are append-only across restarts with size-based rotation, and every emitted alert is journaled with the per-channel delivery status – so "I never received X" is answerable in seconds from records, not reconstruction. Push delivery itself is dual-transport (two independent services), either sufficing for deduplication.

The model-data layer has a complete free fallback: if the primary model API is down for >45 min, a reduced grid is fetched from MET Norway's Locationforecast (no key; the published terms – 20 req/s ceiling, identified User-Agent, If-Modified-Since caching – are honored by construction, with the fallback grid sized to ~5,100 requests/day), and the grid consumer serves it transparently: downstream code cannot tell the source changed, and the one missing variable (CAPE) degrades to the satellite-derived stability product. Satellite data gaps after power or connectivity loss are backfilled up to 24 h from the public archive by a startup catch-up thread, protecting the continuity of the verification dataset. Section 6.10 documents the day all of this was exercised at once.

5 Verification Protocol

5.1 Infrared Hindcast

For each base time t_k, the forecast uses only frames available at t_k and is compared against the observed field at t_k + lead, repeated over all possible base times of a continuous series – 77 frames (13:00Z-01:40Z, 2026-06-03/04), giving N = 62-73 pairs per lead; the S-PROG/ANVIL bake-off used ~90 frames. The baseline is persistence – the honest control in nowcasting, since the atmosphere nearly persists at short leads; a nowcast that cannot beat persistence has a motion field worth nothing. Metrics: MAE of brightness temperature; POD/FAR/CSI for the Tb $\leq$ -55 °C event (CSI ignores correct negatives, the honest choice for rare events); FSS at 2/10/50 km scales to separate position from structure skill. A regrid cache plus threaded base-computation reduced a full verification run from ~24 min to ~1.9 min (12.6x), with results identical to two decimals – performance claims in this project are measured, never asserted.

5.2 Radar Hindcast

Radar verification uses dBZ $\geq$35 as the event and four contrasting events (a 2026 clean advective case, a 2019 clean case at a second radar, a 2018 mixed 108 mm event, a 2019 hail/chaotic case; 21-91 frames each) – small by satellite standards because the public API retains no archive, which is itself a finding about this data regime: forward logging is the only way to build a verification dataset, and the system has logged continuously since 2026-06-08. CSI/FSS results are insensitive to the precipitation threshold (5/15/25 dBZ tested identical), and the method *ranking* is immune to the colorbar bias of Section 4.1 (relative comparisons on the same field), while absolute metrics required reprocessing after the bias fix.

5.3 Convective-Initiation Prototype

The CI diagnostic evaluates Lagrangian cloud-top cooling against a strict physical proxy: a cumulus region (-50 to -20 °C) that advection does not carry into an existing storm but that subsequently crosses the deep-convection threshold. This is an experimental diagnostic, not an alert product.

5.4 Forward Operational Verification

Since 2026-06-08 every alert decision is logged with its full input snapshot (the *pred*), and a separate service labels each record ~2 h later with observed truth: maximum reflectivity within 5 km of the target, lightning within 30 km (with coverage flags), and rain/no-rain. This forward log is the substrate for the calibration program of Section 9; first aggregate statistics appear in Section 6.9.

6 Results

6.1 Infrared Advection Beats Persistence

Pure infrared advection outperformed persistence in MAE at every lead from +10 to +120 min (Figure 3a):

Lead	MAE nowcast	MAE persistence	POD	FAR	CSI
+10 min	1.75	3.32	0.85	0.06	0.81
+30 min	3.70	5.80	0.65	0.11	0.60
+60 min	5.65	7.93	0.43	0.16	0.40
+90 min	7.07	9.26	0.28	0.19	0.26
+120 min	8.24	10.22	0.16	0.25	0.15

A repeat run with 80 frames reproduced MAE to two decimals and POD/CSI within 0.02 – the result is stable, not a sampling accident. The optical-flow field contains real advective information; exact cell placement weakens beyond about one hour.

6.2 FSS Separates Position Skill from Structure Skill

Lead	FSS ~2 km	FSS ~10 km	FSS ~50 km
+10 min	0.87	0.96	0.99
+30 min	0.69	0.81	0.92
+60 min	0.49	0.59	0.76
+90 min	0.35	0.42	0.56
+120 min	0.23	0.28	0.37

With FSS >0.5 as the useful-skill criterion, exact position is useful to roughly +50 min while 50 km structure remains useful to nearly +100 min (Figure 3b). The growing 2 km / 50 km gap quantifies "the mass is right, the position is not" – and explains why a field can look meteorologically plausible while failing a point-alert threshold.

6.3 S-PROG Is the Best Tested Infrared Method

Lead	MAE pure	MAE S-PROG	MAE ANVIL	CSI pure	CSI S-PROG	CSI ANVIL
+10 min	2.01	1.56	1.52	0.76	0.79	0.80
+60 min	7.27	**5.12**	5.54	0.32	**0.39**	0.35
+120 min	10.54	**7.11**	8.05	0.12	**0.20**	0.16

S-PROG and ANVIL beat pure advection in all three metrics at all 12 leads. Between them S-PROG wins CSI at 11/12 leads; ANVIL's multiscale sophistication does not materialize on transformed Tb. The mechanism: S-PROG's per-scale autoregression discards the small-scale noise advection carries forward, and CDF probability matching restores the intensity distribution rather than washing out cold tops. A methodological note we believe generalizes: *a priori* skepticism against S-PROG was wrong, and an *a priori* plausible Lagrangian growth/decay correction (Section 6.4) was harmful – both directions were resolved only by benchmark.

6.4 Raw Lagrangian Growth/Decay Failed

Persisting the non-advective Tb tendency with exponential damping worsened pure advection in MAE and CSI at all 12 leads, for any damping constant (the limit of zero recovery being pure advection itself). Frame-to-frame Tb tendency is too noisy to persist directly. The negative result is retained to prevent reintroduction of a plausible but harmful correction.

6.5 Radar Requires Regime-Adaptive Method Selection

On radar reflectivity no fixed method is robust (Figure 4): S-PROG wins clean advective events (CSI 0.43 vs 0.31 persistence at +9 min), persistence wins the chaotic hail event where S-PROG collapses (0.13 vs 0.03 at +27 min), and the mixed event splits by lead. The recent-skill adaptive gate equals or approaches the best fixed method in all four events without knowing the regime a priori – in the mixed event it is best at every lead – and never collapses. Flow coherence (not speed, not explosiveness) is the discriminating diagnostic. The physical interpretation: infrared cloud tops evolve smoothly across scales (a permanently pro-S-PROG regime), while radar convection switches regimes within a single event.

6.6 Infrared Is Not Rainfall Truth, but Has System-Scale Value

Radar-IR comparison during convection shows the expected physical relationship (stronger dBZ under colder tops), but the best Tb threshold for dBZ $\geq$35 yields low CSI and high FAR at pixel scale: anvils overspread surface rain, and winter shallow precipitation hides under warm tops. Radar is the correct rainfall truth for local alerts; infrared earns its place for system-scale anticipation and the clearing question (no echo + no cold cloud above the target). The two sources also measure *different velocities* (Section 4.8 Layer 2: cloud-top motion ~77 km/h median vs echo motion ~16 km/h): never averaged, the top velocity caps the cell velocity, and their difference is a free estimate of vertical wind shear.

6.7 Pixel-Level Convective Initiation Has Signal but Fails Operationally

The Lagrangian cooling-rate prototype achieved POD 0.83/0.56/0.41 at +10/+30/+60 min under a strict physical definition – state-of-the-art detection levels – with FAR ~ 1.00. A C13-C15 split-window cirrus filter removed 92% of detections but destroyed POD (0.81 to 0.06) while FAR stayed ~1.00: initiation is a ~1:10,000 event at pixel level, and no pixel filter changes that arithmetic. This reproduces the field's consensus from the inside: the operational unit must be the *object* (segmented, tracked cumulus with per-object features and lightning/radar confirmation), not the pixel.

6.8 Operational Latency Improvements

Measured engineering changes with direct forecast-value impact: processing the local radar before all context radars; event-driven alerting on new local grids (with watchdog); truth labeling and regional refresh moved out of the urgent path; alerts shipped with the latest cached map (the new map renders after); GLM consumed from local aggregates in the urgent path. Net effect: a representative fusion run dropped from ~92 s to ~3.1 s, and no alert waits for any product not required for the local decision.

6.9 Forward Operational Verification: First Statistics

First aggregate from the forward log (570 truth-labeled decisions, dominated by a winter stratiform episode; convective-season statistics pending):

Emitted state	N	Rain ($\geq$35 dBZ within 5 km) in ~2 h	Rate
raining	26	23	88%
arriving	19	6	32%
regional	350	61	17%
clear	172	38	22%

Two honest readings. The "raining" state verifies at 88% (the remainder is intermittent weak echo at the truth threshold). The most interesting number is the **22% rain rate under "clear"**: the dominant misses are weak stratiform bands forming *in situ* near the target – there is no front to detect approaching, because nothing approaches. This identified a missing signal (local echo tendency within 10 km) now on the roadmap, and illustrates what forward verification is for: the system's blind spots are measured, not guessed.

6.10 Operational Case Studies

(a) The RLAN episode (2026-06-10/11). Mutating interference wedges on three radars (Figure 5) generated at least one spurious regional alert before the QC chain reached its final form. The episode is reported in full – symptom, failed geometric attempts (including the sub-windowing fix rejected by its regression test), and the physical resolution (Layer 6 vetoes, validated on 80,444 echoes) – because the failed attempts carry as much design information as the fix: there is no single-frame geometric discriminator between interference bands and radially aligned rain bands, and the escape is cross-source physics, not better geometry.

(b) The double-outage day (2026-06-11). Within one day the national radar network's local radar stopped publishing (~8 h, provider side) and the model API became unreachable (server-side

TLS failures confirmed from three network origins). The system degraded as designed: the source monitor pushed impact statements within minutes; the regional layer fell back to satellite rainfall, which correctly tracked a real approaching system; temperature advection switched transparently to the free fallback model source (139 grid points from met.no, fetched within its published terms); and on provider recovery both feeds resumed with recovery notifications, while the radar gap (irrecoverable by design of the provider's 110-minute archive) remained covered by the satellite record for verification purposes. A stress test no synthetic drill would have composed.

(c) Advection versus propagation (2026-06-11). Documented in Section 4.9 and Figures 6-7: the radar tracker and the satellite trackers disagreed about approach, both correctly, and the in-situ pressure tendency arbitrated. The episode hardened the design rule that motion disagreement is information, never grounds for inter-source suppression.

7 Discussion

7.1 Latency Is Part of Lead Time

Nominal forecast lead is not useful lead. A +60 min forecast computed from a product that is already 30 min old at publication is a +30 min decision advantage at best. RADARES treats source latency as part of the model: the local radar dominates nearby urgent decisions because it is fresher and geometrically more precise, even though it sees less far; the regional layer was rebuilt (own composite) when the official product's latency was measured at twice the radar's. For household-scale nowcasting the stakes are concrete: when a storm is 30-60 km away, waiting for a freshly rendered multi-layer map can consume a meaningful fraction of the available action time – so the alert ships with the latest cached map, and the new map renders after.

7.2 Source Hierarchy Beats Averaging

The sources do not estimate one latent variable with different noise: radar measures hydrometeor backscatter, IR measures cloud-top temperature, GLM counts optical transients, RRQPE retrieves rain rate, the pressure sensor integrates the mass field. Averaging them mixes semantics. RADARES uses hierarchical decision logic – the source with the right semantics and geometry answers each question; others qualify confidence or severity.

The same argument applies to *motion*, not just intensity. The advection-versus-propagation case (Sections 4.9, 6.10c) shows two sources measuring genuinely different physical quantities – core advection versus system propagation – reaching opposite conclusions about approach, both correctly. A fusion design that let the higher-resolution source veto the other would have silenced the only signal anticipating the system's expansion toward the target; averaging the vectors would have produced a motion corresponding to no physical process at all. Hierarchical fusion with per-source announcements preserves the disagreement as information, and an independent in-situ predictor (surface pressure tendency) arbitrates at negligible cost.

7.3 The Unit of Prediction Matters

Pixel-level initiation is hopelessly imbalanced; object-level initiation matches the phenomenon (cumulus objects develop, split, merge, glaciate, electrify). The same lesson generalizes within this

system: the approach-gating entity criterion (Section 4.9) is an object-level constraint retrofitted onto a field detector, and the alert lifecycle treats the *announced system*, not the pixel field, as the unit with state (announced, tracked, consolidated, stood down).

7.4 Why the Result Is Transferable

The coordinates and the radar network are local; the design pattern is not:

1. Recover physical fields from public rendered products when raw data are unavailable – and validate the conversion (colorbars, georeferencing) before believing any number derived from it.
2. Measure every source's latency and subtract it from nominal lead.
3. Assign sources by role and semantics; never average across semantics; let motion disagreement live as information.
4. Demand zero real-signal loss from every QC layer, on real-event regression tests that run before each change ships.
5. Use adaptive method choice when field regimes change within the archive.
6. Keep slow context, verification, and rendering out of the urgent alert path.
7. Give alerts a full lifecycle – gating, consolidation, stand-down – because the absence of a message is also a message.

Many countries expose public weather images but not research-ready volumes. RADARES demonstrates a complete scientific workflow in that constraint regime on consumer hardware.

8 Limitations

- Reflectivity is recovered from rendered PNG color classes (5 dBZ quantization), not raw moments; absolute intensities carry that floor of uncertainty.
- The radar verification archive is small (four events) because the public API retains no history; the adaptive-gate result is a regime insight, not a climatological score. Forward logging since 2026-06-08 is the remedy in progress.
- Infrared verification covers one continuous diurnal cycle; seasonal breadth is pending.
- The forward-verification sample (570 decisions) is winter-dominated; the alert thresholds of Section 4.9 (entity size, approach cone, ETA ceiling) are winter-calibrated and will be revisited against summer convection, where a small fast-growing cell may earn its first alert one or two cycles later than ideal.
- Parallax for the visual IR layer uses a per-pixel correction from estimated cloud height; the ACHA-retrieved height is used for detection logic but not yet for the full visual layer.

- The alert product is deterministic and rule-based; calibrated probabilities require the forward archive to mature.
- Pixel-level CI remains non-operational by design (Section 6.7).
- The system is a research prototype, not an official warning authority.

These limitations define the research program: accumulate forward truth across seasons, report reliability and Brier scores, validate the radar gate leave-one-event-out, move CI to object level, and only then train learned models.

9 Future Work

1. **Operational verification at scale**: hundreds of forward-logged cases with truth labels for rain, intensity, lightning, and clearing; reliability diagrams and Brier scores; per-distance and per-source POD/FAR.
2. **Useful-lead reporting**: skill tables with measured source latency subtracted from nominal lead.
3. **Calibrated probabilities** for heavy rain within 30/60 min and lightning within 30 km, replacing binary alerts.
4. **Leave-one-event-out validation of the radar adaptive gate** as the archive grows.
5. **Object-level CI**: segment and track cumulus objects; per-object features (area growth, Tb minimum, Lagrangian cooling, split-window, water vapor, GLM proximity, radar echo, motion coherence); gradient-boosted classification in the ProbSevere spirit [Cintineo2018].
6. **Corridor transit matching**: the same system identified in successive radars along the 400+ km climatological approach corridor (Cordoba-Parana-Pergamino-Ezeiza) yields measured – not extrapolated – regional approach velocities, upgrading regional ETA confidence from low to medium-high. Registration is in production; matching awaits the first real fronts.
7. **Learned models last**: LightGBM or deep nowcasting after, not before, a sufficient local archive and an honest validation protocol.

10 Conclusion

RADARES demonstrates that a scientifically disciplined point-nowcasting and alerting system can be built entirely from public, imperfect weather imagery on consumer hardware. The findings with the widest reach are not tied to Buenos Aires: S-PROG is the best tested deterministic method on geostationary infrared while radar demands regime-adaptive selection; quality control for imagery-derived reflectivity ends at cross-source physics because single-frame geometry cannot separate interference bands from radially aligned rain; motion disagreement between radar and satellite trackers can be the signature of propagation – information to preserve, not error to suppress; the alert lifecycle needs an explicit stand-down because silence after a warning is an unstated forecast; and latency belongs inside the model, not in an engineering appendix.

The strongest scientific value of the system is the explicit coupling of verification, latency, source semantics, quality control, and point-decision design – with negative results reported alongside positive ones, and every operational rule traceable to a measured episode.

Data and Code Availability

The system consumes exclusively public data: SMN/SINARAME radar imagery, NOAA GOES-19 ABI/GLM products from the AWS Open Data registry, NOAA/AWC METAR, the Open-Meteo API, MET Norway Locationforecast, and ERA5. The implementation (Python 3.11; numpy, scipy, netCDF4, pyproj, pySTEPS, matplotlib) resides in a research repository; the target coordinate is anonymized in this manuscript. Verification scripts regenerate every table and figure from archived inputs.

Acknowledgments

This project uses public meteorological imagery and data from Argentina's SMN/SINARAME, NOAA GOES-R series (GOES-19 ABI and GLM), NOAA/AWC, ECMWF ERA5, Open-Meteo, and MET Norway. It builds on the open-source pySTEPS ecosystem for motion estimation, nowcasting, and verification.

References

[Ayzel2020] Ayzel, G., Scheffer, T., Heistermann, M. (2020). RainNet v1.0: a convolutional neural network for radar-based precipitation nowcasting. *Geoscientific Model Development*, 13, 2631-2644. https://doi.org/10.5194/gmd-13-2631-2020

[Bedka2010] Bedka, K., Brunner, J., Dworak, R., Feltz, W., Otkin, J., Greenwald, T. (2010). Objective satellite-based detection of overshooting tops using infrared window channel brightness temperature gradients. *Journal of Applied Meteorology and Climatology*, 49(2), 181-202. https://doi.org/10.1175/2009JAMC2286.1

[Cintineo2018] Cintineo, J. L., Pavolonis, M. J., Sieglaff, J. M., Lindsey, D. T., Cronce, L., Gerth, J., Rodenkirch, B., Brunner, J., Gravelle, C. (2018). The NOAA/CIMSS ProbSevere Model: Incorporation of Total Lightning and Validation. *Weather and Forecasting*, 33(1), 331-345. https://doi.org/10.1175/WAF-D-17-0099.1

[GermannZawadzki2002] Germann, U., Zawadzki, I. (2002). Scale-dependence of the predictability of precipitation from continental radar images. Part I: Description of the methodology. *Monthly Weather Review*, 130(12), 2859-2873. https://doi.org/10.1175/1520-0493(2002)130<2859:SDOTPO>2.0.CO;2

[LucasKanade1981] Lucas, B. D., Kanade, T. (1981). An iterative image registration technique with an application to stereo vision. *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, 674-679.

[MecikalskiBedka2006] Mecikalski, J. R., Bedka, K. M. (2006). Forecasting convective initiation by monitoring the evolution of moving cumulus in daytime GOES imagery. *Monthly Weather Review*, 134, 49-78. https://doi.org/10.1175/MWR3062.1

[Mecikalski2008] Mecikalski, J. R., Bedka, K. M., Paech, S. J., Litten, L. A. (2008). A statistical evaluation of GOES cloud-top properties for nowcasting convective initiation. *Monthly Weather Review*, 136(12), 4899-4914. https://doi.org/10.1175/2008MWR2352.1

[MetNo] Norwegian Meteorological Institute. Locationforecast 2.0 and API Terms of Service. https://api.met.no/ and https://developer.yr.no/doc/TermsOfService/

[NCEIABI] NOAA National Centers for Environmental Information. GOES-R Terrestrial Weather: ABI/GLM. https://www.ncei.noaa.gov/products/goes-terrestrial-weather-abi-glm

[NOAAGLM] NOAA GOES-R Series. Geostationary Lightning Mapper. https://goes-r.noaa.gov/spacesegment/glm.html

[NOAAOpenData] Registry of Open Data on AWS. NOAA GOES 16/17/18/19. https://registry.opendata.aws/noaa-goes/

[NOAAStarBand13] NOAA STAR. GOES ABI Band 13 (10.3 µm clean longwave infrared window) fact sheet. https://www.star.nesdis.noaa.gov/goes/

[OpenMeteo] Open-Meteo. Free weather API. https://open-meteo.com/

[Pulkkinen2019] Pulkkinen, S., Nerini, D., Perez Hortal, A. A., Velasco-Forero, C., Seed, A., Germann, U., Foresti, L. (2019). Pysteps: an open-source Python library for probabilistic precipitation nowcasting (v1.0). *Geoscientific Model Development*, 12, 4185-4219. https://doi.org/10.5194/gmd-12-4185-2019

[Ravuri2021] Ravuri, S., Lenc, K., Willson, M., et al. (2021). Skilful precipitation nowcasting using deep generative models of radar. *Nature*, 597, 672-677. https://doi.org/10.1038/s41586-021-03854-z

[RobertsLean2008] Roberts, N. M., Lean, H. W. (2008). Scale-selective verification of rainfall accumulations from high-resolution forecasts of convective events. *Monthly Weather Review*, 136(1), 78-97. https://doi.org/10.1175/2007MWR2123.1

[Schultz2009] Schultz, C. J., Petersen, W. A., Carey, L. D. (2009). Preliminary development and evaluation of lightning jump algorithms for the real-time detection of severe weather. *Journal of Applied Meteorology and Climatology*, 48(12), 2543-2563. https://doi.org/10.1175/2009JAMC2237.1

[Seed2003] Seed, A. W. (2003). A dynamic and spatial scaling approach to advection forecasting. *Journal of Applied Meteorology*, 42(3), 381-388. https://doi.org/10.1175/1520-0450(2003)042<0381:ADASSA>2.0.CO;2

[Shi2015] Shi, X., Chen, Z., Wang, H., Yeung, D.-Y., Wong, W.-K., Woo, W.-C. (2015). Convolutional LSTM network: a machine learning approach for precipitation nowcasting. *Advances in Neural Information Processing Systems*, 28.

[SINARAME] Servicio Meteorologico Nacional (Argentina). Sistema Nacional de Radares Meteorologicos (SINARAME). https://www.smn.gob.ar/radar

[Zhang2023] Zhang, Y., Long, M., Chen, K., et al. (2023). Skilful nowcasting of extreme precipitation with NowcastNet. *Nature*, 619, 526-532. https://doi.org/10.1038/s41586-023-06184-4

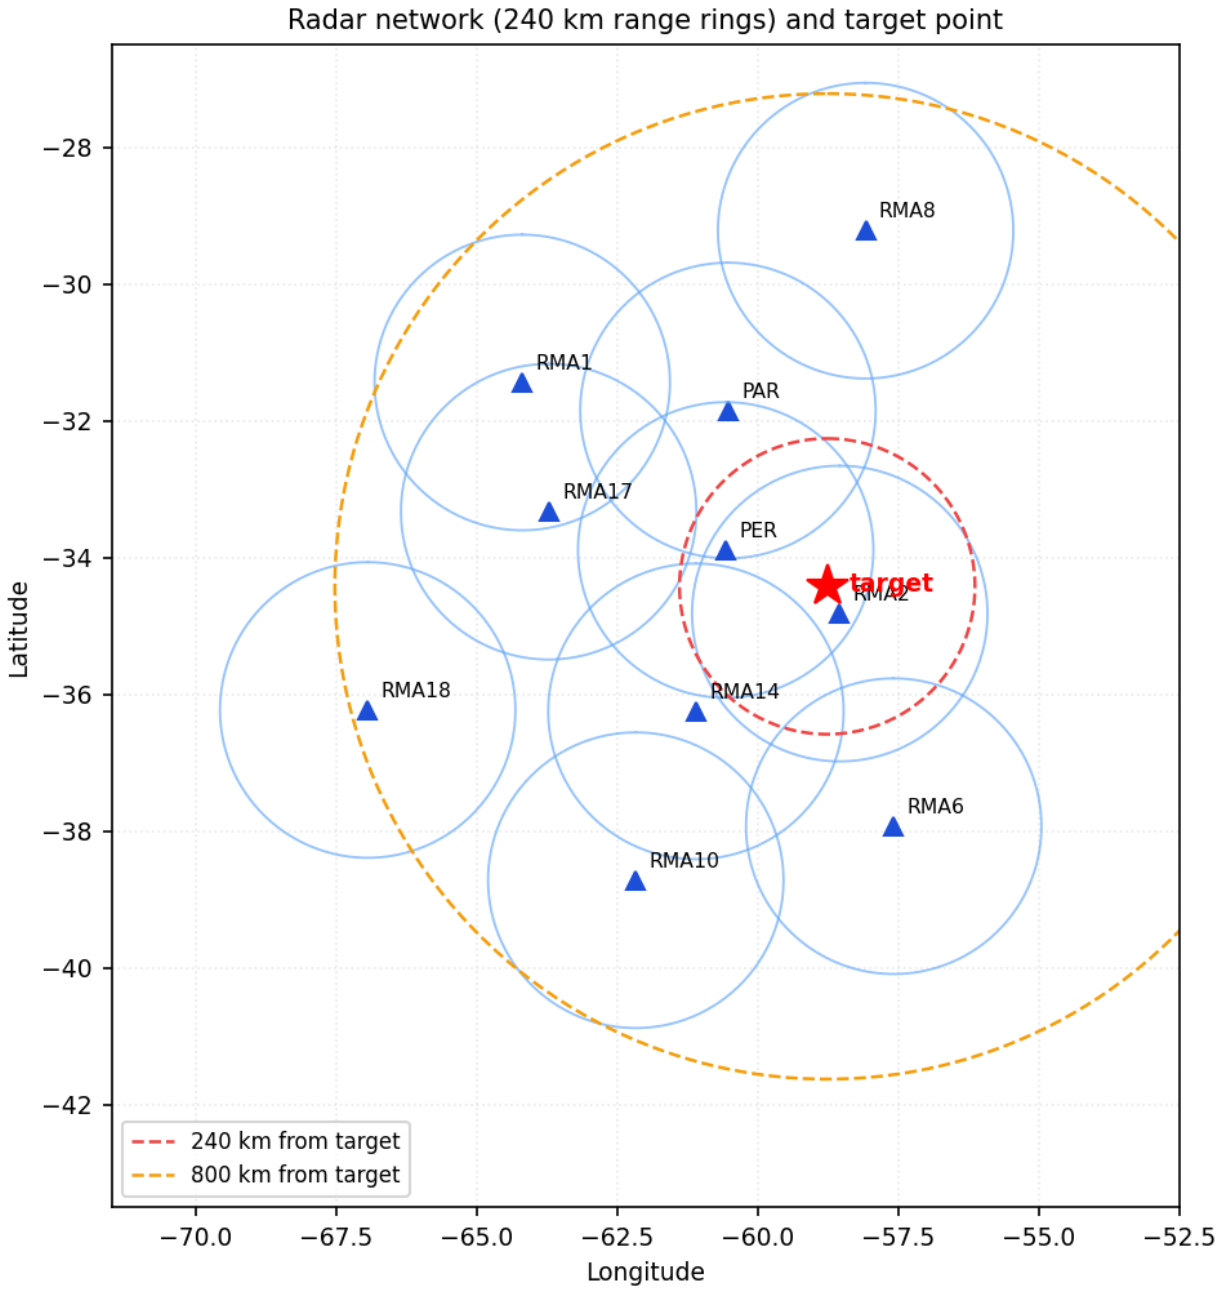

Figure 2: the ten calibrated SINARAME radars with their 240 km range rings, and the target point with 240 km (local/regional boundary) and 800 km reference rings.

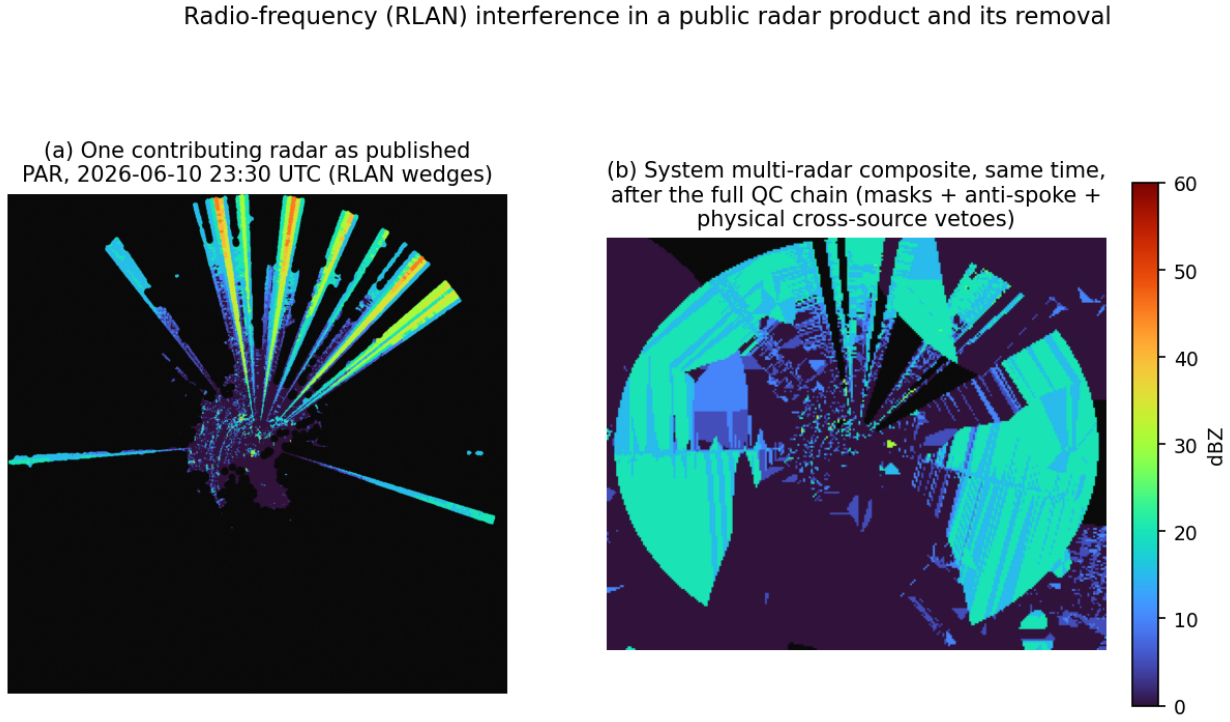

Figure 3: (a) one contributing radar as published by the public portal during the 2026-06-10 RLAN episode – radial interference wedges reach 30-45 dBZ and mutate azimuth between frames; (b) the system's multi-radar composite at the same instant after the full QC chain: the wedges are gone while the genuine shallow-precipitation field survives.

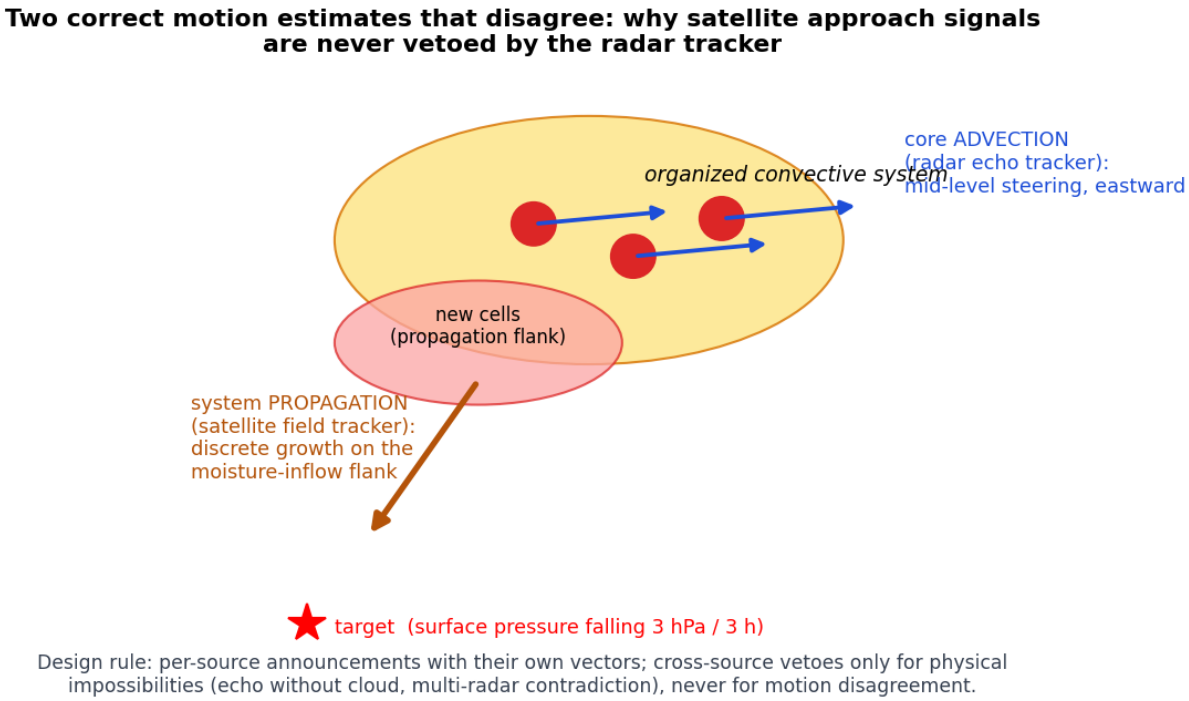

Figure 4: schematic of the documented advection-versus-propagation case. Radar echo tracking (blue) measures core advection with the steering flow, eastward and away from the target; satellite field tracking (orange) measures system propagation by new-cell growth on the moisture-inflow flank, toward the target – confirmed independently by falling surface pressure at the target. Both are correct; vetoing or averaging them would destroy information.

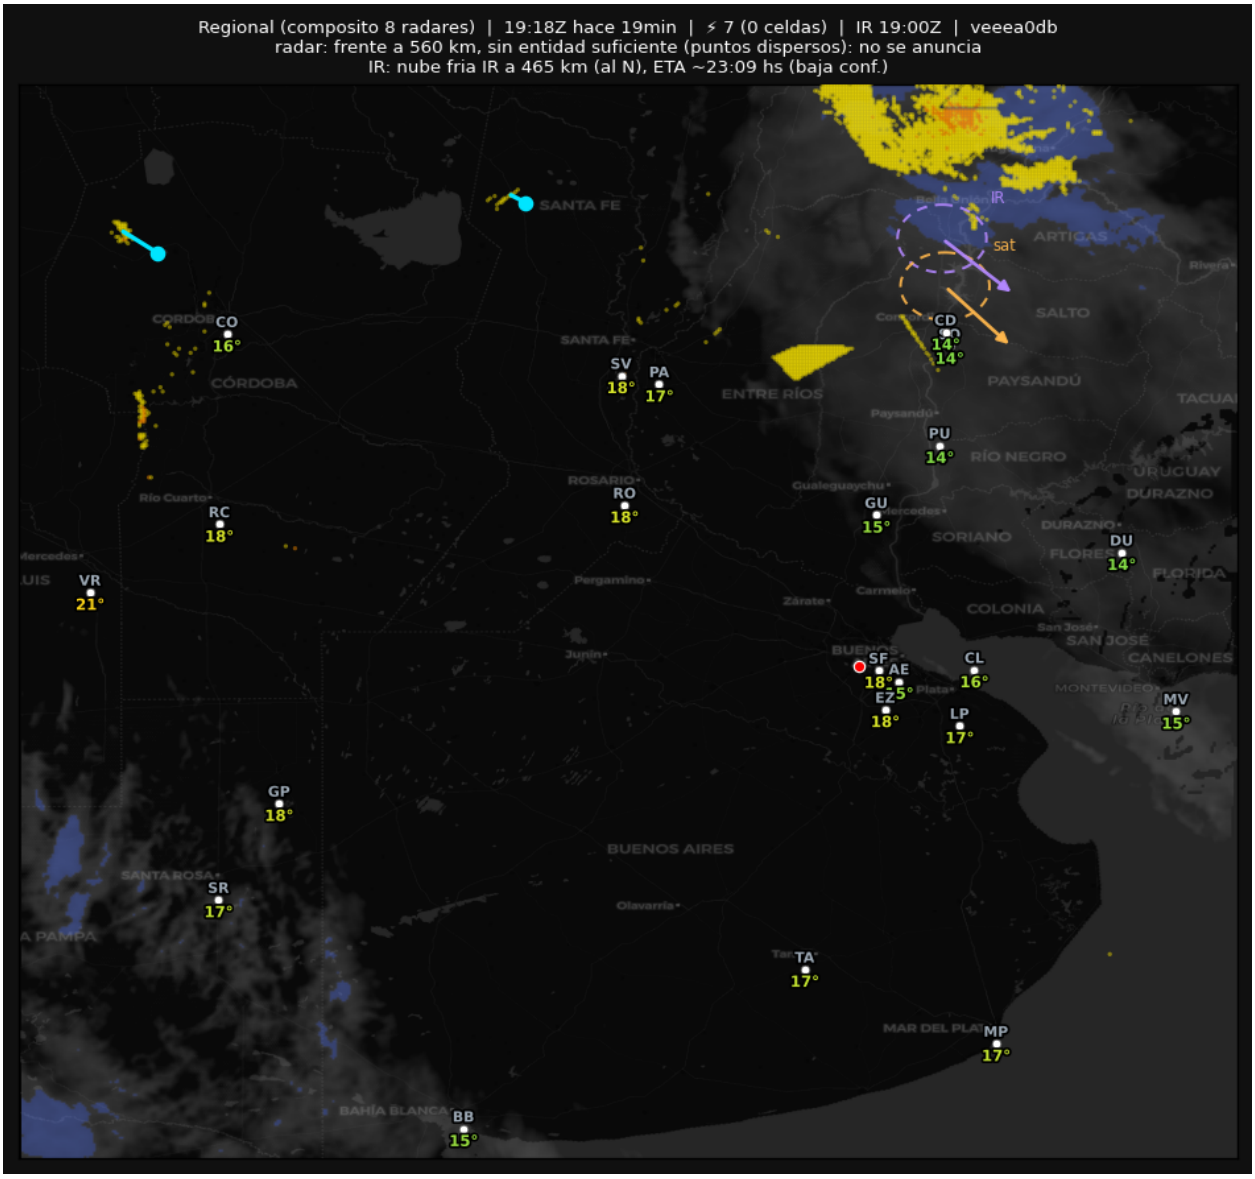

Figure 5: the operational regional map during the documented case. Radar echoes (color), infrared cloud layers (grey/blue), GLM lightning, METAR temperatures, and the announced-object markers: the orange dashed ring and vector labeled "sat" mark the object announced by the satellite rain-rate source on the system's southern flank; the caption line reports the radar detection as "too far in time: not announced" – the same gate as the push channel.

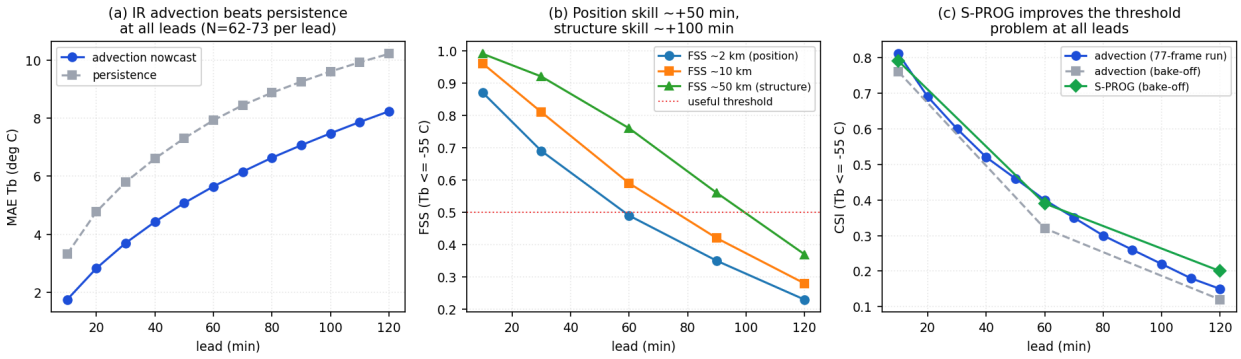

Figure 6: infrared verification. (a) Advection beats persistence in MAE at all leads. (b) FSS by scale: position skill (2 km) is useful to ~+50 min, structure skill (50 km) to ~+100 min – the growing gap is the signature of advection error. (c) The threshold problem: S-PROG improves CSI over pure advection at all leads in the bake-off.

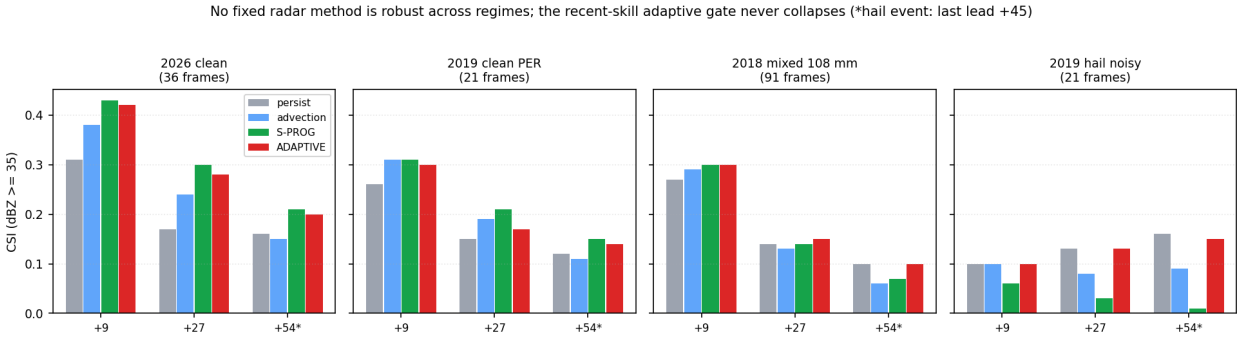

Figure 7: radar nowcast CSI by method and lead across four contrasting events. No fixed method is robust – S-PROG wins clean events and collapses in the hail case where persistence wins – while the recent-skill adaptive gate (red) tracks the best method in every regime.