

Weak-Signal Interpretation for Search-and-Rescue Sensing Pipelines

Preserving Weak Clues, Widening Attention, and Predicate-Bound Promotion

Stephen A. Putman

putmanmodel@pm.me

Version 0.2 — May 2026

Abstract

Search-and-rescue sensing pipelines often begin not with cleanly resolved person-location claims, but with weak, scattered, noisy, and locally ambiguous clues: faint audio reports, weak thermal irregularities, partial tracks, intermittent phone or beacon signals, low-confidence drone-camera cues, witness fragments, terrain disturbance, or animal/search-dog indications. In such cases, the architectural problem is not only detection, but promotion control: how weak findings should influence attention before they are allowed to become stronger search significance.

This note proposes a minimal, inspectable surface: a light state ladder, a bounded clue object, promotion predicates, and compact examples showing how weak clues can widen attention without triggering tasking. It argues that SAR sensing pipelines are a suitable application domain for a stratified weak-signal architecture organized around bounded findings, provisional clue interpretation, short-horizon attention widening, and governed promotion under human/team review.

The proposal does not define a full search-and-rescue command system, autonomous dispatch platform, person-identification system, operational doctrine, or replacement for trained search coordinators. Its narrower contribution is architectural: SAR systems need a disciplined place for weak clues to matter before they harden into tasking decisions, search hypotheses, or ignored noise.

1. Problem Framing

Search-and-rescue operations rarely begin with certainty. More often, they begin with incomplete and uneven clues distributed across terrain, sensors, people, weather, time, and local circumstance. A faint sound may be reported once and then disappear. A phone ping may be weak or stale. A drone image may contain an object-like shape that could be a person, debris, shadow, animal, or terrain feature. A partial footprint may be meaningful, but only in relation to route, weather, substrate, and last-known-position assumptions.

The engineering problem is to preserve and correlate weak, heterogeneous clues while preventing premature operational hardening. Search-and-rescue needs a disciplined middle state: weak clues should be able to widen attention without immediately becoming stronger claims about where a missing person is located, where teams should be sent, or what operational significance has been established.

This bridge paper treats SAR sensing as a weak-signal interpretation problem. The goal is not to replace human judgment, field expertise, incident command, or local protocols. The goal is to define a minimal architecture in which weak clues can be preserved, compared, correlated, and promoted only under declared criteria.

2. Transfer Thesis

The transfer claim is narrow:

Search-and-rescue sensing pipelines appear to be a suitable application domain for a weak-signal architecture in which bounded clues raise attention first, candidate meanings remain provisional, and stronger consequence occurs only through persistence, convergence, rationale preservation, and human/team review.

Assumptions: (a) sensors and reports can be timestamped and geo-referenced with known uncertainty; (b) incident command or the local SAR workflow can accept a provisional attention layer; and (c) human reviewers are available for S4 promotions. If these assumptions fail, the architecture's utility is limited.

This framing does not reject search planning, trained field judgment, GIS tools, drone sensing, dog teams, acoustic sensors, thermal imaging, phone forensics, or conventional SAR practice. It reframes the early-stage handling of ambiguous clues. In many SAR contexts, the first useful evidence is too partial to deserve strong commitment but too important to ignore.

The governing design principle is simple:

Weak search clues should influence attention before they justify stronger operational consequence.

That means a weak phone ping, a faint audio report, an uncertain thermal shape, a partial track, or a witness fragment can widen a search window, trigger corroboration, preserve a candidate area, or request additional inspection without immediately becoming a person-location claim or tasking decision.

The value of the architecture is disciplined attention under uncertainty. It gives the system a place to hold provisional significance while still requiring stronger support for stronger consequence.

3. Stratified Architecture Mapping

The proposed architecture separates SAR clue handling into four stages.

3.1 Bounded Observation and Provisional Interpretation

At the first layer, the system emits narrow findings only. These findings remain local to their source, sensor, report, or observation path. Examples include a faint call-like audio event, a weak

thermal irregularity, a partial footprint, a possible fabric fragment, an intermittent beacon signal, a low-confidence drone image cue, a search-dog indication, or a witness report fragment.

At the second layer, the system forms candidate local meanings from one or more bounded clues. These are not final truths. They are admissible interpretations of a partial pattern. Examples include possible recent passage, possible human presence, possible stale clue, possible environmental artifact, possible false report, possible route deviation, possible sheltering area, possible water drift indicator, or unresolved clue requiring more evidence.

The key requirement is that ambiguity remain visible. The system is not more trustworthy because it hides uncertainty; it is more trustworthy when candidate meanings remain revisable.

3.2 Short-Horizon Attention Widening

At the third layer, weak but informative clues influence short-horizon attention. Short-horizon actions are limited to: (1) schedule a corroborating sensor pass within N minutes; (2) expand the watch radius by M meters; (3) set a decay timer T; or (4) flag for higher review priority. Values for N, M, and T are context-dependent and should be set by local SOPs. None of these actions authorizes team dispatch.

The important constraint is that this layer remains temporary and non-authorizing. It shapes what the system inspects next, but it does not by itself dispatch teams, rewrite durable assumptions, declare location, or authorize stronger operational consequence.

3.3 Governed Promotion / Human-Team Review

Only after persistence, convergence, or stronger support builds should the architecture permit governed promotion. At this stage, the system may allow escalation into a human/team review queue, candidate-area prioritization, resource-routing consideration, or search-plan adjustment.

Promotion checklist (S3 → S4):

- Predicate record: list predicates satisfied, such as persistence, convergence, spatial plausibility, or temporal plausibility.
- Supporting evidence: attach clue IDs and timestamps.
- Uncertainty envelope: include the combined location uncertainty polygon or bounded search area.
- Reviewer token: preserve reviewer identity or role, decision timestamp, and decision outcome.
- Actionable outcome: one of {escalate to tasking review, schedule field inspection, archive as R0, quarantine as Q0}.

Promotion is allowed only when the checklist is complete and a human reviewer signs the reviewer token. S4 makes a case eligible for human operational decision; it does not automatically dispatch teams.

A compact pipeline expresses the structure clearly:

local clue → candidate meaning → bounded attention widening → governed promotion

Or, in the broader stack language:

notice → interpret → shape → govern

4. Minimal Operational Contract

This section keeps the transfer concrete without turning the paper into a full SAR platform specification. The point is not to prescribe one implementation. It is to define a minimal contract surface that makes the architecture inspectable and portable.

4.1 Light State Ladder

A minimal state ladder is enough to distinguish weak-signal attention from stronger operational consequence:

- **S0 — Background:** no meaningful clue state beyond ordinary search context.
- **S1 — Noticed:** a weak clue is recorded as a bounded finding.
- **S2 — Correlating:** the clue persists, recurs, or aligns with another weak partial.
- **S3 — Candidate Area / Candidate Trail:** convergent clues support a provisional search hypothesis.
- **S4 — Promoted for Review / Tasking:** a governed boundary has been crossed; team review or resource adjustment becomes eligible.
- **R0 — Resolved:** the clue is explained, dissipates, loses support, or is superseded by stronger evidence.
- **Q0 — Quarantined:** evidence favors artifact, duplicate, stale report, unreliable source, environmental confound, or sensor error.

Operational semantics:

- **S1 Noticed:** clue stored with provenance and decay timer; no cross-correlation required.
- **S2 Correlating:** automatic cross-checks run; `supporting_clues[]` may be populated.
- **S3 Candidate:** provisional search polygon or bounded candidate area created; short-horizon actions allowed.
- **S4 Promoted:** human review required before any resource allocation or tasking change.
- **R0 Resolved / Q0 Quarantined:** archived with closure reason and retention policy.

The key boundary is simple: S1-S3 are attention and evidence states; S4 belongs to governed review and operational workflow.

4.2 Minimal Clue Object

Minimal clue JSON schema, required fields:

```
{
  "clue_id": "string",
  "source_type": "enum[audio,thermal,visual,phone,beacon,track,witness,dog]",
  "timestamp_utc": "ISO8601 string",
  "location": {"lat": "number", "lon": "number"},
  "location_uncertainty_m": "number",
  "modality": "string",
  "observed_signal": "string",
  "confidence": "number",
  "candidate_state": "enum[S1,S2,S3,S4,R0,Q0]",
  "supporting_clues": ["clue_id"],
  "review_horizon_minutes": "number",
  "decay_rule": "string",
  "promotion_status": "string",
  "rationale": "string"
}
```

Example instance:

```
{
  "clue_id": "c-20260505-001",
  "source_type": "phone",
  "timestamp_utc": "2026-05-05T14:12:00Z",
  "location": {"lat": 42.1234, "lon": -73.5678},
  "location_uncertainty_m": 120,
  "modality": "cell-tower-ping",
  "observed_signal": "intermittent ping",
  "confidence": 0.35,
  "candidate_state": "S1",
  "supporting_clues": [],
  "review_horizon_minutes": 180,
  "decay_rule": "linear-decay",
  "promotion_status": "none",
  "rationale": "initial ping near route edge"
}
```

The point is not to freeze a universal schema. The point is to preserve enough context that weak clues remain reviewable, comparable, and governable.

4.3 Transition Logic

Transition logic should remain explicit. The following predicates are illustrative and must be tuned to local SOPs, terrain, sensor reliability, and incident conditions.

- $S0 \rightarrow S1$: any bounded observation with provenance and timestamp.
- $S1 \rightarrow S2$: the same clue is observed twice, or an independent clue appears within 200 m and 60 minutes.

- S2 → S3: at least two independent modalities fall within the combined uncertainty polygon, or a supporting track and witness report align within 120 minutes.
- S3 → S4: the promotion checklist is complete and combined confidence is ≥ 0.7 , or a human reviewer overrides the threshold.
- S1/S2/S3 → R0: the decay timer expires, or explicit negative confirmation is received.
- S1/S2/S3 → Q0: source reliability falls below threshold, or sensor error is flagged.

This ladder prevents a weak clue from jumping directly into strong consequence. It also prevents the opposite failure: a weak clue can remain live long enough to attract corroboration before it is discarded.

5. Promotion Predicates

Promotion from provisional clue state into stronger search significance should require more than salience alone. Example predicates include:

- persistence across a bounded review window
- convergence across two or more clue types
- spatial consistency with last-known position, route, terrain, or drift model
- temporal plausibility relative to disappearance timeline and movement constraints
- source reliability or independent corroboration
- compatibility with environmental conditions
- preserved rationale package sufficient for later review
- human/team review before stronger tasking or resource adjustment

These are examples, not fixed operational thresholds. The architectural requirement is that promotion be declared, inspectable, and predicate-bound rather than silently triggered by a single weak clue.

Each predicate must be computable from clue fields and logged in the promotion record. For example, convergence means at least two clues with non-overlapping failure modes whose location uncertainty polygons intersect; persistence means the same clue type is observed at least twice separated by at least X minutes. X should be set by local SOPs and incident conditions.

A promotion record should preserve the candidate label, supporting clues, relevant location uncertainty, timing assumptions, confidence regime, promotion predicates satisfied, reviewer or governance token if applicable, and closure path if later revised.

6. Compact Examples

Example 1: Lost Hiker

Worked trace:

- T0: Phone ping c-20260505-001; state=S1, location_uncertainty_m=120, confidence=0.35.
- T+12m: Witness shout c-20260505-002; state=S1, location_uncertainty_m=80, confidence=0.25.
- T+45m: Field team logs track disturbance c-20260505-003; state=S1, location_uncertainty_m=10, confidence=0.60.
- Automatic correlation: the system computes the intersection of uncertainty polygons. The intersection is non-empty. supporting_clues[] is updated for each record. All three clues move to S2.
- S2 → S3: predicate satisfied. At least two modalities, phone and track, fall within a bounded review window and overlapping polygon. The system creates candidate polygon P1 and sets review_horizon_minutes=180.
- Promotion attempt: promotion checklist assembled. Combined confidence is computed from phone and track and weighted by modality: 0.62, below the illustrative 0.70 threshold. No automatic S4 is allowed. The system schedules a corroborating drone pass as a short-horizon action.
- If drone confirms a thermal anomaly within P1: combined confidence rises above threshold; promotion checklist is complete; human reviewer signs reviewer token; case moves to S4 and candidate area becomes eligible for tasking review.

Minimal promotion record:

```
{
  "candidate_id": "P1",
  "candidate_state": "S4",
  "supporting_clues": ["c-20260505-001", "c-20260505-003", "c-20260505-004"],
  "predicates_satisfied": [
    "convergence",
    "spatial_plausibility",
    "temporal_plausibility"
  ],
  "combined_confidence": 0.74,
  "uncertainty_envelope": "polygon:P1",
  "reviewer_token": {
    "reviewer_role": "SAR_coordinator",
    "decision_timestamp_utc": "2026-05-05T15:25:00Z",
    "decision": "promote_for_tasking_review"
  },
  "rationale": "phone ping, track disturbance, and drone thermal cue supported P1"
}
```

This trace shows the core boundary: weak convergence can create a candidate area, but stronger operational consequence requires checklist completion and human review.

Example 2: Disaster Rubble Search

A sensor pass detects a faint sound-like event inside a debris field. Thermal imaging shows a weak irregularity, but the scene contains heat-retaining materials and unstable surfaces. A void-map analysis suggests a plausible survivable pocket nearby.

The sound-like event and thermal irregularity are bounded findings. The interpretation layer holds competing possibilities: possible survivor presence, structural noise, heat artifact, or unresolved anomaly. Attention widening may request another acoustic check, a different viewing angle, or closer review by trained personnel. Promotion requires convergence, persistence, and human/team review before stronger operational consequence.

The architecture helps prevent both false certainty and premature dismissal.

Example 3: Shoreline or Water Search

A floating item is observed downstream from the last-known area. A weak visual cue from drone footage appears nearby, but current and wind conditions complicate interpretation. A witness also reports seeing movement earlier along the shoreline.

The item, visual cue, and witness report remain separate bounded clues. The system may form a provisional drift-compatible candidate area if timing and current models align. Attention may widen downstream inspection or request another drone pass. Stronger search-plan consequence requires governed promotion and review, especially because water movement can rapidly create misleading spatial associations.

The value is not automatic conclusion. The value is disciplined preservation of a possible clue cluster.

7. Why This Matters

Search-and-rescue sensing is high consequence precisely because weak clues may matter before they are cleanly resolved. The cost of ignoring them can be severe. The cost of overcommitting to them can also be severe.

This architecture reduces two operational errors: premature tasking from single noisy clues, and premature dismissal of corroborating weak signals.

8. Scope and Non-Goals

This paper does not propose:

- a full search-and-rescue command system
- an autonomous dispatch platform
- a person-identification system

- a certified emergency-response framework
- a replacement for trained SAR personnel
- an operational doctrine
- a guarantee of rescue success
- a universal SAR data schema
- automated authority to task teams or allocate resources

The claim is narrower. SAR sensing pipelines need a disciplined place for weak clues to matter before they harden into tasking decisions, durable search hypotheses, or ignored noise.

9. Conclusion

Search-and-rescue sensing pipelines operate under uncertainty, consequence, and incomplete evidence. The earliest useful signals are often not resolved person-location claims, but weak clue clusters that deserve wider attention without yet deserving operational certainty.

A stratified weak-signal architecture provides a disciplined response. Bounded clues are preserved. Candidate meanings remain provisional. Attention widens under short-horizon significance. Stronger search consequence requires governed promotion, rationale, and human/team review.

Concretely: implement the S-ladder, adopt the minimal clue schema, log promotion records, and require a human reviewer for S4. The next step would be a short pilot integration with one sensor modality and a documented set of promotion thresholds; operational readiness should not be claimed until such validation exists.