Cessation Signature Search: A Statistical Framework for Detecting Technosignature Loss in Multi-Epoch Survey Archives

Arnar Ragnarsson

Independent Researcher

contact@arreport.com

7 April 2026

Abstract

I propose a methodological reorientation of technosignature search programmes. The standard SETI paradigm monitors for the emergence of anomalous signals against background. I argue that the complementary question—has the composite background itself changed in ways inconsistent with known astrophysical processes?—is not the explicit target of any systematic programme: existing multi-epoch surveys are generally optimised for resolved-source variability rather than temporal stability of the unresolved composite residual—the aggregate flux remaining after subtraction of catalogued point sources and best-available diffuse emission models. Exploring it constitutes an independent search channel whose sensitivity properties are distinct from resolved-source emergence detection. I term this approach the Cessation Signature Search (CeSS). For the purposes of public communication, the same programme is also designated Negative SETI, or N-SETI. CeSS's core proposition is that technosignatures may be more detectable as perturbations in statistical fields than as discrete signals—and that the perturbations caused by their cessation have not been searched for. Its primary object is the detection of statistically significant changes in the aggregate signal background, especially in the unresolved composite residual, while treating discrete source disappearance as the more tractable end of the same detection spectrum. Relevant archival data already exists to support initial CeSS pilot studies: decades of radio and infrared survey data whose composite statistical properties have never been characterised for temporal stability. Signal loss or transformation may have various causes, each implying a different class of historical pattern in the signal record, though some are only partially discriminating and may overlap observationally. I describe the observational and statistical framework, define the detection spectrum from composite deviation to discrete dropout, identify natural sterilising events as a control channel, and state the falsifiability conditions. Portions of this work were generated with the assistance of artificial intelligence systems, including models from GPT, Claude, and Gemini; full details are provided in the AI Use and Disclosure Statement.

Keywords: technosignatures; SETI; radio surveys; multi-epoch observations; composite residuals; Fermi paradox

---

1. Introduction

The search for extraterrestrial intelligence has, since its inception, been organised around a single epistemic posture: look for something new. Surveys are designed to detect anomalous signals—emissions that should not exist under purely natural astrophysical processes—appearing against a background that is treated as a static baseline to be subtracted. This framing has shaped instrument design, pipeline architecture, and target selection across six decades of organised search (Tarter, 2001). I argue that this framing is systematically incomplete in a specific and addressable way.

The observable universe does not present us with a clean background plus a few candidate signals. It presents us with a composite: the superposition of electromagnetic contributions from an enormous number of sources, the vast majority unresolved, uncatalogued, and individually undetectable. Standard SETI pipelines attempt to isolate anomalous components from this composite. They are not designed to ask whether the composite itself has changed—whether its statistical character at a given sky position has shifted in ways that cannot be explained by known astrophysical processes. Existing multi-epoch variability programmes are optimised for resolved-source flux changes; the stability of the unresolved composite residual as a technosignature observable has not been systematically explored.

The approach has structural precedents in other fields that similarly extract information from absence, indirection, or statistical residuals. Exoplanet transit photometry detects planets by measuring the decrease in stellar flux as they occult their host star: absence of light as positive signal. Gravitational wave background searches operate on the composite strain from an unresolved population of sources, inferring properties of that population from statistical properties of the aggregate rather than from individual detections. Cosmic microwave background analysis extracts cosmological parameters from statistical deviations in a field that is, at any individual point, dominated by foreground noise. In each case, the information of interest lives not in individual resolved sources but in the statistical character of the composite. CeSS proposes to apply the same inferential structure to the technosignature problem.

The shift may be subtle. A technological civilisation that never produced a signal strong enough to be individually detected might nonetheless have contributed measurably to the aggregate background in its sky region. If that civilisation ceased—for any reason—the composite would change. The change would not necessarily look like a signal disappearing. It would look like noise changing character: a slight shift in spectral composition, flux density, or variability statistics that is inconsistent with the known processes governing that region. Detecting such a shift requires a specific kind of instrument: a statistical monitor of background stability. This is the primary object of CeSS.

The detection spectrum runs from the subtle—a statistically anomalous shift in composite background properties at a sky position, not attributable to any single source—to the dramatic: the disappearance of a resolved, catalogued emission source with no natural explanation. Both ends of this spectrum are scientifically valuable, and neither is currently the explicit target of any search programme. The dramatic end is rarer and easier to interpret; the subtle end is more demanding. A complete CeSS programme addresses the full spectrum.

A critical observation motivates beginning this programme now: relevant archival data already exists that may support initial CeSS analyses. Decades of radio survey observations and infrared sky maps have been accumulated—and processed through pipelines optimised to find new point sources, but not to characterise composite background stability over time. The assertion that SETI has detected zero technosignatures means zero confirmed, individually identified signals. It says nothing about the statistical properties of the aggregate background across epochs, because those properties have never been computed for this purpose. The data exists; the question has not been asked of it.

The distinction from existing programmes should be stated. Radio variability surveys such as VAST (Murphy et al., 2021) and ThunderKAT flag flux changes of individually resolved sources; radio transient pipelines flag the emergence of new sources above a sensitivity threshold; and cross-epoch image differencing searches for sources that have appeared or vanished between maps—all variants of single-source framing. Confusion-noise analyses characterise aggregate flux from unresolved populations but are not optimised for temporal stability as an observable. Radio stacking recovers emission from known-position populations below the single-source detection floor, but requires prior positional knowledge. All of these programmes are therefore united by two constraints CeSS does not share: their detection statistic is the flux change of a resolved individual source rather than a change in the statistical properties of the unresolved composite residual, and sources below the individual detection limit are excluded from variability analysis entirely. None asks primarily whether the statistical character of that unresolved composite residual has changed between epochs in a way that cannot be explained by known astrophysical processes. That is the question CeSS asks.

The causes of signal loss are numerous; the following represent principal categories but do not constitute an exhaustive taxonomy: civilisational self-termination through war, ecological collapse, or technological accident; elimination by another species (Ćirković, 2018; Liu, 2008); voluntary cessation of detectable emissions; transition to non-detectable emission modes (Sandberg et al., 2017); stellar engineering (Dyson, 1960); natural astrophysical sterilisation; and civilisational transformation, relocation, or propagation (Freitas, 1980). These categories imply different classes of historical pattern in the signal record, some of them only partially discriminating and some observationally overlapping. CeSS does not assume that every case can be cleanly classified; rather, it proposes that the structure of the cessation record may constrain the relative plausibility of competing explanations where sufficient data exists.

To illustrate what such patterns look like in practice, Earth provides a useful if humbling example. Viewed from a distant vantage point, Earth's aggregate electromagnetic signature has undergone profound transformation over the past century: the rise of AM radio broadcasting in the 1920s introduced strong modulated carrier signals in the medium-wave band; the subsequent emergence of FM radio, television, and military radar expanded the profile across HF, VHF, and microwave frequencies; the transition to digital broadcasting from the 1990s onward altered the spectral texture of those emissions substantially, shifting from analogue carrier signatures toward noise-like spread-spectrum profiles; and satellite communications have progressively redirected significant transmitter power away from the planetary surface and into directed orbital geometries.

This is signal transformation rather than cessation—a structured change in the character, spectral distribution, and directionality of an aggregate contribution, without any discontinuity in the civilisation's existence. Whether Earth's own emissions would be detectable from even the nearest stellar system is a separate question: at a distance of 4.24 light-years (Proxima Centauri), Earth's aggregate radio leakage flux density is of order $10^{-35}\ \mathrm{W\,m^{-2}\,Hz^{-1}}$

under the assumptions of an isotropically radiated total power of approximately 10^{10} W distributed across a bandwidth of 10^{10} Hz. Earth is not currently detectable at interstellar distances through its leakage alone. But a civilisation emitting at levels substantially higher than Earth's—or observed at closer range, or integrated over sufficiently long baselines—would present precisely this class of transformation signature in the historical record. The CeSS methodology is designed to detect it.

The Fermi paradox—the absence of confirmed technosignatures despite the apparent abundance of habitable environments—has generated a large family of proposed resolutions (Hart, 1975; Webb, 2002). Each, if operative, would be expected to map differently onto the cessation record, though not always in a uniquely discriminating way.

2. The Detection Spectrum

CeSS operates across a continuous spectrum of detection difficulty. At one end is the discrete dropout: a resolved, catalogued emission source—a radio emitter, an anomalous infrared excess, a directed laser signal—that was present in the signal record and subsequently absent, without a natural astrophysical explanation. This is the most interpretively tractable form of cessation event: the source was known, its cessation is directly observable, and causal analysis can proceed against a well-defined prior. It is also the rarest case, since it requires the prior existence of a source strong and distinct enough to have been individually identified.

At the other end is the composite deviation: a statistical change in the aggregate background properties of a sky region that is inconsistent with known astrophysical evolution but is not attributable to any single resolved source. This case arises when a contributing civilisation was never individually detectable—its emission was below the detection threshold for single-source identification but contributed measurably to the aggregate. Composite deviation is a signal-processing and statistical inference problem rather than a source-identification problem. It requires stable, long-baseline models of the background composite against which temporal deviations can be measured. Between these extremes lies a range of intermediate cases: partially resolved sources whose contribution to the composite is uncertain; sources detectable in stacked or co-added data but not in individual observations; and regions where the composite has changed in ways that are consistent with the cessation of one of several candidate contributors. The statistical methodology must be designed to operate across this range, with appropriate confidence levels assigned to detections at each point.

I define a cessation candidate as a sky region satisfying all four of the following criteria.

(a) The composite residual $R(\alpha, \delta, \nu, t)$ has changed between at least two epochs by an amount exceeding the predicted variance from known astrophysical processes at the relevant timescale, as quantified against the baseline model.

(b) The change is consistent with the removal or transformation of a structured contribution and is inconsistent with a transient or single-epoch instrumental effect.

(c) The change cannot be attributed to any catalogued astrophysical event within the positional and temporal uncertainty bounds.

(d) The change is corroborated by contemporaneous data in at least one additional frequency band, or the absence of such corroborating data is explicitly documented as a limitation of the candidate record.

A 'structured contribution' is one exhibiting one or more of the following physically motivated properties: narrowband spectral components inconsistent with thermal broadening at local stellar temperatures; anomalous infrared excess at effective temperatures inconsistent with known stellar photosphere models; temporal modulation with non-random statistics (e.g. periodicity or non-Poissonian variability); or spatial coherence inconsistent with diffuse foreground models at the survey beam scale. It should be noted that characterised instrumental drift patterns and scintillation histories are not merely confounds to be excluded; when incorporated into the background model, they constitute additional constraints that sharpen the boundary between natural variance and anomalous deviation, improving the model's discriminating power over successive epochs.

The statistical validation framework for cessation candidates is analogous to that used in radio transient and variable source surveys (Fender & Bell, 2011; Murphy et al., 2021), with one critical difference in direction: transient pipelines flag increases above threshold; cessation pipelines flag decreases below baseline. Both require stable background models, multi-epoch confirmation, and false positive rejection through multi-band consistency. The underlying statistical methodology largely exists in the transient literature; it requires reorientation. It should be noted that detection sensitivity decreases toward the composite end of the detection spectrum, and interpretation correspondingly becomes more model-dependent; the framework's strongest claims are at the discrete dropout end, where individual sources can be localised and causal attribution is more tractable.

A clarification on the epistemological structure of this methodology is warranted. Criterion (a) above requires that a deviation exceed the expected variance from known astrophysical processes. This does not require a perfect model of the background. It requires that the best available model—constructed from the current understanding of galactic emission, source evolution, and instrumental behaviour—be used to generate a predicted variance, and that the observed deviation be evaluated against that prediction. The residual is then a statement about what current models cannot account for, not a claim of certainty about what it is. This is the standard epistemological structure of residual-anomaly science: CMB analyses proceed under incomplete foreground models; exoplanet surveys operate under imperfect stellar variability models; gravitational wave background searches work against incompletely characterised noise floors. In each case, rigour comes not from perfect knowledge of the background but from the honest quantification of its uncertainty. CeSS operates under the same constraint and the same standard.

3. Natural Sterilisation Events as a Control Channel

Several classes of high-energy astrophysical events produce ionising radiation sufficient to sterilise planetary biospheres at interstellar distances (Gehrels et al., 2003; Melott & Thomas, 2011). Because their occurrence, timing, and sky positions are in principle independently observable, they provide a unique control channel: cessation candidates spatially and temporally coincident with a known sterilising event of sufficient fluence are candidates for natural extinction, while those in the residual population require civilisational explanations. If real composite anomalies are found coincident with independently observable sterilising events, this supports a natural-extinction interpretation of the cessation event, but does not by itself establish that the pre-cessation contribution was technological. The sterilising event provides a plausible natural mechanism for the loss of the anomalous contribution; the anomaly itself remains subject to the same interpretive uncertainty as any other second-level residual. Such a coincident population may therefore help classify a subset of cessation candidates under plausible natural-extinction mechanisms, while not by itself constituting confirmation of past civilisational activity or a measure of civilisational prevalence.

A practical limitation of this cross-reference is the incompleteness of the historical sterilisation record: systematic gamma-ray burst catalogues begin with BATSE (1991) and the Fermi GBM (2008), and the magnetar flare catalogue is similarly shallow in historical depth. For cessation candidates whose emission epochs predate the instrumental record, the natural-extinction cross-reference fails, and those candidates must be treated as unclassifiable at this level rather than as residual civilisational events. It should be noted that sterilisation radius calculations carry significant uncertainty from atmospheric shielding models, planetary obliquity and orbital geometry, and habitability assumptions; the affected volumes cited below should be treated as order-of-magnitude constraints rather than precise classification boundaries, and this uncertainty propagates into the confidence with which any individual cessation candidate can be assigned to the natural-extinction category.

4. Observational Methodology

The null and alternative hypotheses of CeSS may be stated at the outset.

- H_0: all multi-epoch changes in the composite residual $R(\alpha, \delta, \nu, t)$ are fully accounted for by instrumental systematics, calibration drift, and known astrophysical variability processes.
- H_1: some subset of those changes requires the addition, removal, or transformation of a structured non-natural contribution—one whose spectral, temporal, or spatial properties are inconsistent with known astrophysical emission classes.

CeSS is the programme of testing H_1 against H_0 on archival and prospective multi-epoch sky data, using the composite residual as the test statistic and multi-band cross-confirmation as the primary rejection criterion against instrumental confounds.

The composite residual $R(\alpha, \delta, \nu, t)$ at sky position (α, δ), frequency ν, and epoch t is defined as the total measured flux density after subtraction of all catalogued point sources and best-available models of diffuse galactic and extragalactic emission:

$$R(\alpha, \delta, \nu, t) \;=\; S_{\text{obs}}(\alpha, \delta, \nu, t) \;-\; S_{\text{cat}}(\alpha, \delta, \nu, t) \;-\; S_{\text{diff}}(\alpha, \delta, \nu, t). \tag{1}$$

CeSS monitors R for temporal changes inconsistent with known astrophysical variance. The core methodological inversion of CeSS is as follows. Standard emergence pipelines ask: is there a component in the current signal that exceeds background? Cessation pipelines ask: has the background itself changed between epochs in a way that cannot be explained by known astrophysical processes? This requires reorienting the primary data quality criterion from point-source sensitivity to composite background stability.

Concretely, the pipeline operates as follows. For a given sky region and frequency band, a background model is constructed from all available observations over the baseline period. This model characterises the mean flux, spectral composition, spatial structure, and variability statistics of the composite. Subsequent observations are then tested against this model: does the current composite deviate from the baseline model by more than expected from known astrophysical processes at the relevant timescale? A statistically significant deviation—quantified as a likelihood ratio against a null model of natural background evolution—

constitutes a cessation candidate requiring multi-band confirmation.

A worked illustration. Consider a sky region whose radio background at 1.4 GHz has been characterised over a ten-year baseline. The composite flux includes contributions from resolved sources (catalogued and subtracted), diffuse galactic emission (modelled), and a residual that is the aggregate of unresolved contributions. If this residual decreases between epochs by an amount that is statistically inconsistent with known variability processes—calibrated against control regions without candidate civilisational contributions—and if the decrease is corroborated by a contemporaneous change in the infrared background of the same region, this constitutes a composite deviation candidate. The multi-band requirement is the primary rejection criterion against single-band instrumental effects, scintillation, and propagation artefacts.

The expected false positive rate for composite deviation candidates under realistic foreground conditions—the number of spurious detections per sky patch per epoch—is not derived here; it is instrument- and site-specific, and its quantification is a key deliverable for any implementation programme. The multi-band cross-confirmation requirement is expected to suppress most single-band artefacts, but a sky-patch-level false positive budget must be established empirically and is explicitly deferred to implementation work. As a structural bounding argument: if the per-band, per-sky-patch-epoch false positive probability is p_1 (determined by the calibration floor and the chosen detection threshold), then multi-band confirmation across k independent bands suppresses the composite false positive rate to approximately p_1^k for uncorrelated systematics, or to $p_1 \cdot \rho$ for correlated systematics where ρ is the cross-band correlation coefficient—in either case, denoting the resulting composite false positive rate as p_{fp}. For a survey of N_{patch} sky patches over N_{ep} epochs, the expected number of spurious composite-deviation candidates is $N_{\text{patch}} \times N_{\text{ep}} \times p_{\text{fp}}$. The multi-band requirement is therefore not merely a confirmation step but the primary lever on the false positive budget.

Throughout this framework, an epoch is defined as a single observing campaign of defined duration at a sky position—its length chosen to be short relative to the cessation timescale of interest but long relative to the natural variability timescale being modelled. The minimum useful temporal baseline is set by the requirement to distinguish cessation from natural variability. For radio and infrared bands, baselines of years to decades suffice to characterise stellar variability and galactic background fluctuations; longer baselines increase sensitivity to slower processes and reduce the false positive rate for composite deviation candidates.

Several specific archives are immediate candidates for a first-pass CeSS reanalysis. In the radio, NVSS (1.4 GHz, full northern sky, 1990s epoch) provides a deep baseline for comparison against modern surveys; VLASS (2–4 GHz, full northern sky, three-epoch cadence from 2017) provides internal epoch-differencing and NVSS cross-epoch comparison; TGSS ADR1 (150 MHz, $\sim$37,000 deg^2) offers a low-frequency complement; the MeerKAT archive (900 MHz–1.7 GHz, ongoing from 2018) provides southern-sky cadence. In the infrared, WISE and NEOWISE (3.4–22 μm, multi-epoch from 2010) enable contemporaneous cross-band confirmation. Breakthrough Listen and the Allen Telescope Array contribute targeted deep observations for selected sky regions. Each of these datasets carries its own calibration history, epoch cadence, and sensitivity floor; archival sufficiency for a CeSS reanalysis is therefore a dataset-by-dataset empirical question requiring harmonisation assessment, not a general property of the radio and infrared archive.

4.1 Detection Limits: Physical, Technological, and Algorithmic

The detectability of Earth-level signals at interstellar distances raises a question that bears directly on the practical scope of CeSS: are the relevant limits set by physical law, by the current technological level of the listening civilisation, or by the absence of appropriate algorithms? The answer matters because it determines whether the composite deviation end of the detection spectrum is a permanently closed door or an open engineering problem. The answer is the latter.

The physical floor. The fundamental sensitivity limit of a radio receiver is set by the radiometer equation:

$$\sigma_S = \frac{\text{SEFD}}{\sqrt{\Delta\nu \cdot \tau}}, \tag{2}$$

where SEFD is the System Equivalent Flux Density, $\Delta\nu$ is the observed bandwidth, and τ is integration time. This is not a ceiling: sensitivity improves without bound as collecting area increases, bandwidth widens, and integration time grows. The quantum noise limit at 1.4 GHz corresponds to a noise temperature $T_Q = h\nu/k_B \approx 0.07\,\text{K}$. Current best cryogenic receivers at that frequency operate at system temperatures of 3–10 K—roughly 1.5 to 2 orders of magnitude above the quantum floor. This gap is not obviously closed by any simple first-principles limit, though reducing it in practice remains a substantial systems-engineering challenge.

The technological gap. The gap between current capability and the quantum limit is an engineering problem. Josephson parametric amplifiers (JPAs), already deployed in quantum computing hardware, achieve noise temperatures within a factor of 2–3 of the quantum limit at microwave frequencies. Adapting them for radio telescope front-ends—an active engineering programme—would reduce system noise temperatures from 5–10 K to 0.1–0.5 K, a 10–$100\times$ improvement in noise power and 3–$10\times$ improvement in flux sensitivity. A globally coordinated, phase-coherent interferometric network—extending the Event Horizon Telescope model to encompass SKA, MeerKAT, FAST, VLA, and ngVLA—would synthesise baselines of planetary diameter, reaching microarcsecond-scale angular resolution at short centimetre wavelengths, with resolution degrading to the milliarcsecond regime at longer centimetre wavelengths. These capabilities are in the engineering pipeline; none require new physics.

The algorithmic gap. Current radio survey pipelines use largely frequentist, threshold-based detection optimised for individual point sources. They are not designed to extract information from signals that are permanently below the individual detection threshold but present in the statistical properties of the aggregate. Several fields have already solved structurally analogous problems, and some of their methods may be adaptable here. Anomaly detection in large time-domain survey datasets has been demonstrated at scale in the optical domain (Malanchev et al., 2021); extending such methods to the composite residual problem in the radio and infrared is a tractable algorithmic step.

Pulsar timing arrays (NANOGrav, EPTA, PPTA) detect the nanohertz gravitational wave background by correlating timing residuals across tens of millisecond pulsars over baselines of 15 or more years. No single observation contains the signal; it exists only in cross-correlated residuals. The 2023 NANOGrav report of evidence for a gravitational-wave background (Agazie et al., 2023; Siemens et al., 2013) at strain amplitude $h_c \sim 10^{-15}$—a signal invisible in any individual measurement—demonstrates the power of long-baseline statistical integration, and the general principle carries a direct implication for CeSS: beginning systematic background monitoring now has compounding returns that cannot be recovered by later investment.

Whether a spatial-covariance analogue of the Hellings–Downs structure can be constructed for CeSS in the spectral domain is an open question requiring explicit derivation before the analogy can be considered operational. The physical basis for such an analogue is motivated: several signal classes produce physically constrained spectral covariance signatures—the cosmic watering hole near 1.4–1.7 GHz, waste-heat Planck spectra, nuclear-weapon broadband pulse profiles, stellar engineering optical-decline/infrared-rise covariance—each defining a template $C(\nu_1, \nu_2)$ that is calculable from physics rather than cultural assumptions.

CMB power spectrum analyses reconstruct cosmological parameters from maps whose individual pixels are dominated by instrumental noise and foreground contamination; the cosmological signal is never detectable at any single pointing. Hierarchical Bayesian models that simultaneously fit source populations, instrumental noise, and propagation effects—approaching the Cramér–Rao bound—recover 2–5$\times$ more signal than classical thresholding at equivalent false positive rates.

Radio stacking already pushes effective detection below nominal survey thresholds: co-adding image cutouts at the positions of optically identified galaxies detects radio emission from populations 5–10$\times$ below the single-source noise floor. CeSS instead requires a field-level statistical analogue of stacking: aggregating information across sky regions without assuming known source positions, and characterising composite background statistics rather than flux at catalogued locations.

For single-beam analysis of a resolved candidate source, the effective detection threshold after N_{ep} independent epochs is $\sigma_{\text{threshold}} = \sigma_{\text{beam}}/\sqrt{N_{\text{ep}}}$. With $\sigma_{\text{beam}} \approx 10^{-32}\ \mathrm{W\,m^{-2}\,Hz^{-1}}$ (SKA1-Mid at 1.4 GHz, one-hour integration, 1 GHz bandwidth; Braun et al. 2019) and $N_{\text{ep}} = 100$ epochs over a decadal programme, $\sigma_{\text{threshold}} \approx 10^{-33}\ \mathrm{W\,m^{-2}\,Hz^{-1}}$. A 100$\times$ Earth emitter at 10 parsecs contributes $8 \times 10^{-35}\ \mathrm{W\,m^{-2}\,Hz^{-1}}$, roughly one order of magnitude below this per-beam threshold.

A critical limitation of both noise floor estimates must be stated explicitly. The calculations above assume thermally uncorrelated noise across beams and epochs. In practice, the dominant floor for composite background stability is not thermal noise but calibration and sky-model systematics: gain drifts between observing sessions, beam shape variation across epochs, antenna position errors, and ionospheric and tropospheric phase corruption are correlated across beams on characteristic angular and temporal scales. Correlated noise does not average down as $1/\sqrt{N_b}$. The practical composite stability floor is therefore set not by thermal noise calculations but by the calibration fidelity of the instrument and the accuracy of the sky model—quantities that are instrument- and site-specific. The multi-band requirement therefore functions as a partial mitigation of the correlated systematics floor, not merely as signal confirmation.

Net assessment. The detection of Earth-level signals at parsec distances is not currently feasible with point-source methods, and will likely not become so within the next decade under existing programmes. Detection of composite deviations from stronger-than-Earth emitters—the more realistic target for the near term—is within reach in terms of thermal noise floors using existing archive data combined with algorithms that are straightforward extensions of deployed techniques in gravitational wave and CMB science, pending characterisation of the calibration systematics floor identified above. The composite deviation end of the detection spectrum is not a permanent boundary set by physics. It is an algorithmic and engineering frontier, and the relevant tools are either already in hand or clearly in view.

4.2 Minimum Implementation Pipeline

The following pipeline operationalises the test of H_1 against H_0 defined in Section 4. The stages are logically sequential: each presupposes the outputs of the preceding stage, and interpretive analysis of surviving candidates applies only to those that have completed Stage 6 in full. Parameter values—thresholds, baseline lengths, confirmation windows—are implementation-specific and must be derived from control-field calibration rather than assumed from first principles.

Stage 1 — Inputs. Multi-epoch survey data covering a target sky field at one or more frequency bands, comprising a minimum of $N_{\text{ep}} \geq 5$ independent epochs separated by intervals sufficient to characterise the dominant natural variability timescale of the background. Required ancillary data: a point-source catalogue for the field complete to the instrument sensitivity limit; a diffuse emission model at the relevant frequency and epoch; and per-epoch calibration records including gain solutions, beam models, and data quality flags.

Stage 2 — Residual construction. For each sky position (α, δ), frequency ν, and epoch t, compute $R = S_{\text{obs}} - S_{\text{cat}} - S_{\text{diff}}$ using calibration solutions specific to that epoch independently. Epochs with degraded calibration quality are excluded from the baseline model but retained in the candidate record with explicit uncertainty flags. The result is a data cube R indexed by sky position, frequency, and epoch.

Stage 3 — Baseline model. Using all epochs within the baseline period established in Stage 1, characterise R at each (α, δ, ν): the time-mean residual $\langle R \rangle$, the per-epoch variance σ^2_{nat}, and the spatial covariance structure across the field. Variance and covariance estimates are calibrated against a set of control fields—sky regions matched to the target field in source density and diffuse emission level, processed through the same pipeline to establish the empirical null distribution of the test statistic.

Stage 4 — Test statistic. For each candidate epoch t_c and sky patch $\mathit{\Omega}$, compute the likelihood ratio $\mathit{\Lambda}(\mathit{\Omega}, \nu, t_c) = P(R \mid H_1)/P(R \mid H_0)$, where H_0 is the baseline model and H_1 is the baseline model plus a free perturbation term representing a change in the contribution from the sky patch. The detection threshold $\mathit{\Lambda}^*$ is set from the control-field null distribution to achieve a target per-field-epoch false positive rate p_{fp}.

Stage 5 — Candidate flagging. A sky patch $\mathit{\Omega}$ at epoch t_c is flagged as a first-level candidate if: (a) $\mathit{\Lambda}(\mathit{\Omega}, \nu, t_c) > \mathit{\Lambda}^*$ in at least one band ν; (b) the deviation is persistent, reproduced above threshold in at least one additional independent epoch within a confirmation window $\mathit{\Delta}T_{\mathrm{confirm}}$; and (c) the deviation is absent from calibration control fields observed in the same session.

Stage 6 — Veto steps. First-level candidates undergo the following sequential tests in order of increasing computational cost.

(i) *Catalogue cross-match*: reject if spatially coincident within positional uncertainty with a catalogued transient, AGN, variable star, pulsed emitter, or other known astrophysical variable class.

(ii) *Calibration and propagation check*: reject if the spatial pattern of the deviation correlates with antenna-based gain errors in the calibration solutions, or if its temporal profile is consistent with interstellar scintillation at the relevant frequency and Galactic latitude.

(iii) *Multi-band confirmation*: retain only if the deviation is corroborated in at least one additional frequency band within a contemporaneous window $\mathit{\Delta}T_{\mathrm{band}}$. Candidates lacking contemporaneous multi-band data are retained with an explicitly downgraded confidence designation; absence of data is not equivalent to absence of signal.

(iv) *Sterilisation event cross-reference*: test for spatial and temporal coincidence with catalogued high-energy sterilising events within the positional and temporal uncertainty bounds on sterilisation radius and propagation delay (Section 3). Coincident candidates are reclassified as natural-extinction candidates rather than removed from the record; such reclassification is itself informative. Where a candidate's anomalous status has been established through the preceding veto steps, coincidence with a sterilising event supports a natural-extinction interpretation of the cessation event, but does not by itself establish that the pre-cessation contribution was technological. The steriliser provides a plausible natural mechanism for the loss of the anomalous contribution; the anomaly itself remains subject to the same interpretive uncertainty as any other second-level residual. This reclassified population is therefore relevant to the programme's classification of cessation mechanisms, but should not be treated on its own as a measure of civilisational prevalence.

Candidates surviving all four veto steps are designated second-level residuals and are eligible for interpretive hypothesis classification.

Stage 7 — Outputs. Outputs are listed in their order of derivation from the preceding pipeline stages; output (d), derived from the complete pipeline, is therefore listed last.

(a) A per-field-epoch table of composite residuals with baseline model parameters and calibration quality indicators.

(b) A ranked first-level candidate list with Λ values, frequency bands, and persistence flags.

(c) A second-level residual catalogue with confidence designations and full veto outcomes.

(d) An empirically determined false positive rate derived from control-field analysis. Output (d) is the primary calibration product of any CeSS implementation: without it, the confidence designations in output (c) are uninterpretable, and no candidate claim can be evaluated.

5. Falsifiability Conditions

The CeSS framework generates the following falsifiability conditions. The methodology itself is subject to empirical validation through its performance on known natural cessation events—astronomical sources confirmed to have gone dark between survey epochs for astrophysical reasons—which constitute control cases that any CeSS pipeline must correctly assign to the natural-process exclusion category before civilisational interpretation is attempted.

(i) *No composite deviations or discrete dropouts detected.* Consistent with either the rarity of technological civilisations, the electromagnetic undetectability of cessation events at achievable sensitivities, insufficient baseline length, or the sensitivity of the initial implementation being insufficient at the emission levels present—as bounded by the analysis in Section 4.1. Does not falsify the methodology but constrains the prior on civilisational prevalence and emission levels.

(ii) *All cessation candidates correlated with sterilising events.* This indicates that natural astrophysical extinction may dominate the accessible signal-loss record. Where candidates' anomalous status has been established through the Stage 6 veto sequence, coincidence with sterilising events strengthens a natural-extinction interpretation of the cessation event, but does not by itself establish that the pre-cessation contribution was technological. Such a result would therefore be informative about the distribution of candidate cessation mechanisms in the accessible record, while leaving broader civilisational-fate hypotheses largely unconstrained.

(iii) *Residual events spatially independent and temporally abrupt, without clustering or propagating structure.* Consistent with self-termination (Hart, 1975; Webb, 2002) or a diffuse predatory regime (Ćirković, 2018; Liu, 2008). Where self-termination proceeds through violent internal mechanisms, the discrete dropout may be preceded by a brief anomalous emission burst in the historical signal record—an observational precursor that, where present, favours self-termination over external elimination or gradual cessation modes.

(iv) *Residual events spatially clustered, or exhibiting a propagating directional structure advancing from a single origin.* Consistent with a territorial predatory regime or a first-mover elimination scenario (Berezin, 2018), and distinguishable from outcome (iii) by the spatial pattern alone.

(v) *Residual events with gradual temporal profiles, or preceded by optical-decline/infrared-rise transitions.* Consistent with voluntary cessation of detectable emissions (Sandberg et al., 2017) or stellar engineering proceeding through progressive energy capture (Dyson, 1960; Wright et al., 2014). The gradual-profile discriminator is most reliable for transitions occurring over timescales substantially longer than the monitoring baseline.

(vi) *Cessation at sky position A followed by correlated emergence at position B, with delay consistent with the A–B light-travel time or sub-relativistic travel.* Consistent with civilisational relocation or von Neumann propagation (Freitas, 1980); identifiable only if emergence-detection and cessation-detection programmes are operated in parallel and their results cross-referenced.

Outcomes (i) through (vi) are not mutually exclusive across sub-populations of the cessation record. A mixed record is the expected case, and some observational signatures may remain only partially discriminating even with substantial data. The framework can nevertheless, in principle, constrain the relative plausibility of different causal categories—self-termination, predatory suppression, voluntary cessation, stellar engineering, and relocation or propagation—provided sample size and observational resolution are sufficient.

6. Conclusion

I have proposed the Cessation Signature Search (CeSS): the systematic monitoring of the signal record for loss rather than emergence. Its primary object is not the dramatic disappearance of a known source but the detection of statistical changes in the composite background—shifts in the aggregate of potentially millions of unresolved contributors, most never individually identified, one or more of which may have ceased. Dramatic discrete dropouts are one end of this detection spectrum; statistical composite deviation is the other. Current search programmes are not explicitly optimised for either.

The framework is grounded in an important empirical observation: relevant archival data already exists that may support initial CeSS analyses, though archival sufficiency is a dataset-by-dataset question. Decades of radio and infrared survey archives contain epoch-by-epoch background records that have never been processed with composite stability as the question. The familiar assertion that SETI has yielded zero results means zero confirmed individual sources above detection threshold—a claim about the emergence pipeline, not about the stability of the composite background it was never designed to monitor.

The framework is strengthened by the existence of independently observable sterilising events. By cross-referencing cessation candidates against the record of gamma-ray bursts, magnetar flares, hypernova radiation, and pulsar exposure histories, a subset of candidates may be classified under plausible natural-extinction mechanisms, while the remaining residual population provides a conceptual framework for comparing how different proposed resolutions to the Fermi paradox (Hart, 1975; Webb, 2002) might map onto the cessation record. The falsifiable predictions of the framework differ across hypotheses, though some are only partially discriminating and may overlap observationally: self-termination; predatory models (Ćirković, 2018; Liu, 2008); civilisational transformation and stellar engineering (Dyson, 1960; Wright et al., 2014); relocation or propagation (Freitas, 1980). These discriminators are unavailable to emergence-based programmes and are not recoverable from the present-epoch silence alone.

There is a final argument that does not depend on any of the above. Cessation detection requires a baseline. Every year of systematic background monitoring that does not occur is a year of baseline that cannot be recovered. Beginning now—with existing archives, existing infrastructure, and new algorithmic tools—compounds future detection sensitivity in a way that later investment cannot replicate. The question is not only what the signal record contains, but how much of it we have been watching. The universe may be full of endings. We have not been looking for them.

Data Availability Statement

No new observational data were generated in support of this research. This work presents a methodological framework. Application of the framework to archival datasets would require access to the radio and infrared surveys cited in Section 4; those datasets are publicly available through the respective survey data archives.

Conflict of Interest Statement

The author declares no conflict of interest.

AI Use and Disclosure Statement

This work was produced with extensive assistance from artificial intelligence (AI) systems across conceptual development, drafting, editing, and analytical support. The author utilized multiple models from GPT, Claude, and Gemini, accessed during March–April 2026.

All AI-generated content was reviewed, evaluated, and curated by the author. The author assumes full responsibility for the accuracy, integrity, and originality of the work, including all interpretations and conclusions presented. No AI system is credited as an author.

References

Agazie, G., et al. 2023, The Astrophysical Journal Letters, 951, L8.

Berezin, A. 2018, arXiv:1803.08425 [preprint].

Braun, R., et al. 2019, Anticipated SKA1 Science Performance, SKA-TEL-SKO-0000818. Square Kilometre Array Organisation.

Ćirković, M. M. 2018, *The Great Silence: Science and Philosophy of Fermi's Paradox*. Oxford University Press.

Dyson, F. J. 1960, Science, 131, 1667.

Fender, R. & Bell, M. 2011, Bulletin of the Astronomical Society of India, 39, 315.

Freitas, R. A. 1980, Journal of the British Interplanetary Society, 33, 251.

Gehrels, N., et al. 2003, Astrophysical Journal, 585, 1169.

Hart, M. H. 1975, Quarterly Journal of the Royal Astronomical Society, 16, 128.

Liu, C. 2008, *The Dark Forest* [Hēiàn Sēnlín]. Chongqing Publishing Group. [English translation: Liu Cixin, Tor Books, 2015.]

Malanchev, K. L., et al. 2021, Monthly Notices of the Royal Astronomical Society, 502, 5147.

Melott, A. L. & Thomas, B. C. 2011, Astrobiology, 11, 343.

Murphy, T., et al. 2021, Publications of the Astronomical Society of Australia, 38, e054.

Sandberg, A., Armstrong, S., & Ćirković, M. M. 2017, Journal of the British Interplanetary Society, 69, 406.

Siemens, X., Ellis, J., Jenet, F., & Romano, J. D. 2013, Classical and Quantum Gravity, 30, 224015.

Tarter, J. 2001, Annual Review of Astronomy and Astrophysics, 39, 511.

Webb, S. 2002, *If the Universe Is Teeming with Aliens… Where Is Everybody?* Copernicus Books.

Wright, J. T., et al. 2014, Astrophysical Journal, 792, 26.